

NAACL-HLT 2021

**The 2021 Conference  
of the North American Chapter  
of the Association for Computational Linguistics:  
Human Language Technologies**

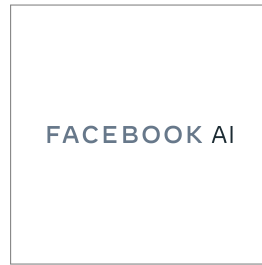
**Proceedings of the Conference**

June 6 - 11, 2021



Diamond Sponsors:



Platinum Sponsors:



Gold Sponsors:



Silver Sponsors:



Bronze Sponsors:



Diversity and Inclusion Champions:



©2021 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](mailto:acl@aclweb.org)

ISBN 978-1-954085-46-6

## Message from the General Chair

It is my pleasure to welcome you to the virtual NAACL-HLT 2021 conference! Although conditions did not allow us to hold the conference in Mexico City as originally planned, we hope our rich program of presentations, poster sessions, discussions, and social events will enable fruitful scientific exchange and increase our connectedness as a community. The increased affordability could also allow us to welcome new members that would not be able to attend a physical conference.

In this message I would first like to highlight a few initiatives and sessions at NAACL 2021 and then acknowledge the many people on our organizing committee and those of prior conferences that were critical to making it happen.

Following suggestions by Dan Jurafsky and members of the NAACL executive board, this year we followed a more well-defined process for ethics reviews<sup>1</sup>, proposed by the ethics co-chairs Emily Bender and Karën Fort, and refined and implemented in coordination with the program co-chairs and the demo and industry track chairs. Key changes were allowing additional space in submissions to discuss ethical considerations and establishing a category of papers accepted conditionally on addressing ethical concerns together with a timeline and process for an additional stage of review of re-submissions. We are grateful to Emily and Karën for setting these foundations and hope our community will continue to improve its process of education about and review of the ethical implications of our research.

Another initiative this year that Graham Neubig suggested, Luciana Benotti, Tamar Solario, Smaranda Muresan, and other members of the NAACL Exec significantly contributed to, and Pranav A and the rest of the D&I committee strongly advocated for and worked hard on was the D&I Grant Initiative.<sup>2</sup> With the help of our sponsors, we were able to waive registration and membership fees for authors from underrepresented developing countries, as well as many others that could not have attended otherwise.

Also new this year was a software package developed by the publication chairs Ryan Cotterell, Steven Bethard, Yichao Zhou, Iz Beltagy, and Tanmoy Chakraborty to automatically check and report formatting violations in an easy-to-understand way. This significant contribution to the infrastructure used by the ACL community went above and beyond the duties of publication chairs for a single conference.

The NAACL virtual conference will host 6 tutorials, 17 system demonstrations, 39 industry track papers, 499 main conference and CL/TACL papers, 6 plenary invited talks, 2 panels, 22 workshops, and a large set of social and thematic gatherings. This was made possible by the hard work of the many members of our organizing committee:

- Anna Rumshisky, Luke Zettlemoyer, and Dilek Hakkani-Tur, our program co-chairs, have contributed the most by leading the selection of the scientific content for the main conference.
- Priscilla Rasmussen arranged our transition to a virtual presence and provided guidance on nearly every aspect of the organization.
- Industry Track Chairs (Owen Rambow, Yunyao Li, and Young-Bum Kim), who advocated for the inclusion of this track and led the selection of 39 papers and additional invited talks and panels.
- Demonstration Track Chairs (Avi Sil and Victoria Lin), who organized the selection of 17 system demonstrations.
- Workshop Chairs (Bhavana Dalvi, Mamoru Komachi, and Michel Galley) who led an efficient and organized process for the workshops despite the uncertainty of the conference format.

---

<sup>1</sup><https://2021.naacl.org/ethics/faq/>

<sup>2</sup><https://2021.naacl.org/blog/dei-grants/>

- Tutorial Chairs (Greg Kondrak, Kalina Bontcheva, and Dan Gillick) who selected and coordinated the presentation of six tutorials.
- Student Research Workshop Chairs (Esin Durmus, Nelson Liu, Vivek Gupta) and Faculty Advisors (Nanyun Peng and Yu Su) who selected 22 research papers and thesis proposals.
- Ethics Chairs (Emily Bender and Karën Fort), who substantially improved the process for ethics review and education.
- Publication Chairs (Ryan Cotterell, Steven Bethard, Yichao Zhou, Iz Beltagy, and Tanmoy Chakraborty) who improved the publication infrastructure through the effort described above.
- Diversity and Inclusion Chairs (Pranav A, Samira Shaikh, Pat Verga, Murathan Kurfali, Khyati Mahajan, and Prathyusha Jwalapuram) and Social Chairs (Luca Soldaini and Sabine Weber) who established the D&I grant initiative, took steps to strengthen the presence of affinity groups at NAACL, coordinated the organization of many socials, mentoring events, and topical discussions, ensured improved accessibility, and distributed financial support to community members in need.
- Publicity Chairs (Sarah Wiegrefe, Enrico Santus, Peng Qi, and Danqi Chen) who made it possible for the program co-chairs and general chair to not have to check Twitter regularly, enabled efficient communication, and initiated a creative way for members to be introduced to the community via a PeopleOfNLProc Blog.
- Volunteer Chair (Hao Cheng) who took the important responsibility of coordinating the work of more than a hundred volunteers.
- Virtual Infrastructure Committee (Deepak Ramachandran, Mauricio Mazuecos, Martín Villalba) for stepping up to secure the foundations of the virtual conference and Advisors (Jan-Christoph Klie, Hao Fang, and Gisela Vallejo) for taking time to point us in the right direction.
- Website Chairs (Ice Pasupat and Iulia Turc) who posted information to the website extremely quickly and in beautiful arrangement.
- Volunteers: More than a hundred volunteers that will help lead the live sessions and ensure information on the website is correct.

I am also grateful to Bonnie Webber, general chair of EMNLP, and Anna Rogers for helping me set expectations on the difficulty of organizing a virtual conference and major challenges to watch out for, and Donia Scott, Horacio Saggion, and Leo Wanner for sharing their experience with Underline, our virtual conference provider, with us. I would also like to acknowledge Sol Rosenberg and Daniel Luise from Underline for arranging everything on a short timeline.

Colin Cherry, David Yarowsky, and other members of the NAACL exec provided valuable advice at multiple decision points.

We are, as always, extremely grateful to our sponsors, listed on the previous page.

Finally, I would like to thank all authors of papers, invited talks, and panels, area chairs, and reviewers, and the volunteers organizing and chairing sessions, and all attendees and readers of this volume for engaging with the content and the community.

Kristina Toutanova  
 NAACL 2021 General Chair  
 June 2021

## Message from the Program Chairs

Welcome to the 2021 Annual Conference of the North American Association for Computational Linguistics! NAACL-HLT 2021 is a completely virtual conference, in response to the COVID-19 pandemic, which will still be felt around the world for some time to come. We are grateful for the support and contributions of the entire NAACL 2021 team. Without them, it would not have been possible to organize an exciting and memorable event during very stressful global times.

We largely adopted the best practices of other recent virtual conferences, while modifying the format somewhat to treat all papers equally, as described in more detail below. Our paper review process followed the recent trend of a hierarchical organization, with senior area chairs (SACs) that organized coherent research tracks and area chairs (ACs) who shepherded smaller batches of papers within each track. We made an effort to balance the tracks. Recent NLP conferences have had many tracks that received well over 200 submissions, making them mini conferences of their own. To make the senior area chairs jobs more manageable, we split the machine learning track by areas (*Classification and Structured Prediction Models* and *Language Modeling and Sequence to Sequence Models*) and separated *Machine Translation* from *Multilingual*. We also wrote a brief guide to authors to help them decide which track was most appropriate for their work.<sup>3</sup> Overall, we felt that this did make the organization more manageable and that other tracks, including NLP applications, could possibly benefit from further splitting in future conferences. Otherwise, we followed recent traditions in track selection, including keeping some of the smaller, more recent additions (e.g. Green NLP and Ethics).

We also had a special theme for the conference, which we called “New Challenges in NLP: Tasks, Methods, Positions.” This theme was selected to recognize that we have made significant progress in NLP over the last five years, and that the community could benefit from thinking about the new problems and upcoming challenges we should focus on next. Despite the general applicability of the unsupervised pre-training/fine-tuning paradigm, many problems are still very challenging for current models. At the same time, given the recent progress, there are likely broad new classes of problems that can now be studied for the first time. What tasks or capabilities should we focus on next? What new classes of models should we be investigating? We envisioned papers falling into this theme including (but not limited to) (1) empirical and dataset papers that propose new challenges that bring us closer to human-level language understanding and generation, and (2) position papers framing an important direction or highlighting an understudied research problem.

We recruited reviewers through a centralized process, designed to minimize workload for senior area chairs (SACs) without sacrificing review quality. We collected a list of likely qualified reviewers based on the reviewer and author pool of other recent NLP conferences. These candidates were invited to sign up to review, and were required to fill out a profile that allowed us to better assess their potential area fit and experience levels. We were fortunate enough to have enough volunteers to not need everyone, and were able to bias the final selection towards more senior reviewers, although many junior and first time reviewers remained in the pool. The final program included 54 SACs, 267 ACs, and 1941 committee members. We greatly appreciate the incredible amount of work they all did, and also thank all of the volunteers who were not selected in the end.

NAACL-HLT 2021 received 1797 submissions—a record for our conference! We accepted 477 papers, including 350 long and 127 short, for an overall acceptance rate of 26%. The acceptance rate for long papers was higher than short papers (28% vs. 23%), although this gap was smaller than in other recent conferences. From the accepted papers, and based on the nominations from SACs and review by the best paper committee, we selected best papers in the long and short paper categories, as well as a small number of outstanding papers in each category. NAACL-HLT 2021 will also feature 18 papers that were published at *Transactions of the Association for Computational Linguistics* (TACL) and 4 papers from

---

<sup>3</sup><https://2021.naacl.org/calls/area-descriptions/>

the journal of *Computational Linguistics* (CL).

Our virtual format includes both interactive talk sessions and posters in Gather.Town, following recent successful online events. A 12-minute video will also be available for long papers and a 7-minute video for short papers. We made a significant change by removing the distinction between oral and poster papers, and instead treat every paper the same. The papers are grouped into sessions, where each session has two parts. The first 50 min will be a live Zoom-like session with oral pitches for each paper, including 5 minutes for long papers and 3 minutes for short papers, followed by 3 minutes of Q&A for each paper. The remainder of the interactive session will happen in Gather.Town, with poster presentations for each paper. The goal is to allow the audience to engage into an in-depth discussion with the authors if desired. This is new model that likely will not scale to very large conference, but will hopefully allow for rich interactions at the conference while only requiring each author to be present for a single session. The sessions were also arranged across time zones to, as much as possible, be in normal waking hours for the main contact authors of each paper. Finally, another highlight of our program is keynote talks from Dhruv Batra, Shakir Mohamed, Hinrich Schütze, and Tamar Solorio. These four excellent speakers were chosen to provide exciting and thought-provoking perspectives from both within and outside of our the core NLP.

NAACL 2021 is truly a community-run effort. We want to second all of Kristina's acknowledgements in the General Chair statement. We have also had the pleasure to work with these folks, and appreciate all of their help. We would additionally like to thank:

- The Senior Area Chairs, who were incredibly organized and responsive, in every step from assigning reviewers to making final decisions. They were a crucial support for the very large area chair and reviewer pools.
- The Area Chairs who led paper review discussions and wrote meta-reviews.
- The primary reviewers and secondary reviewers who provided valuable feedback to the authors. Special thanks to those who stepped in at the last minute to serve as emergency reviewers.
- Our distinguished Best Paper Committee: Isabelle Augenstein, Marco Baroni, Jacob Eisenstein, Hanna Hajishirzi, Omer Levy, Jessy Li, Yang Liu, Chris Quirk, Barbara Di Eugenio, and Bonnie Webber.
- The authors who submitted their work to NAACL-HLT 2021. Although we could not accept many of the strong papers, we know that most of it will end up at other amazing venues, and hope we were able to provide some useful feedback.
- TACL editors-in-chief Mark Johnson, Ani Nenkova, and Brian Roark, TACL Editorial Assistant Cindy Robinson, and CL Editor-in-Chief Hwee Tou Ng for coordinating TACL and CL presentations with us.
- The Program co-Chairs of NAACL 2019, Christy Doran and Tamar Solorio; of ACL 2020, Joyce Chai, Natalie Schluter, and Joel Tetreault; and of EMNLP 2020, Trevor Cohn, Yulan He, and Yang Liu. You were all amazingly available and generous with your time to answer the very many questions we had about how to run a successful program committee.
- And last but not least, our General Chair Kristina Toutanova. She has done an excellent job of steering a large ship in very challenging times!

We sincerely appreciate your help, and hope you will enjoy the NAACL-HLT 2021 conference!

Dilek Hakkani-Tur, Anna Rumshisky, and Luke Zettlemoyer  
NAACL 2021 Program Committee Co-Chairs  
June 2021



# Organizing Committee

## General Chair:

Kristina Toutanova, Google Research

## Program Chairs:

Anna Rumshisky, University of Massachusetts Lowell

Luke Zettlemoyer, University of Washington and Facebook AI Research

Dilek Hakkani-Tur, Amazon Alexa AI

## Local Chair:

Priscilla Rasmussen, Association for Computational Linguistics

## Diversity & Inclusion Senior Chairs:

Pranav A, Miro AI

Samira Shaikh, University of North Carolina - Charlotte

Pat Verga, Google Research

## Diversity & Inclusion Student Chairs:

Murathan Kurfali, Stockholm University

Khyati Mahajan, University of North Carolina - Charlotte

Prathyusha Jwalapuram, Nanyang Technological University

## Diversity & Inclusion Social Chairs:

Luca Soldaini, Amazon Alexa AI

Sabine Weber, University of Edinburgh

## Demonstration Chairs:

Avi Sil, IBM Research AI

Victoria Lin, Facebook AI

## Ethics Chairs:

Emily Bender, University of Washington

Karën Fort, Sorbonne Université

## Industry Track Chairs:

Owen Rambow, Stony Brook University

Yunyao Li, IBM Research

Young-Bum Kim, Amazon

## Publication Chairs:

Iz Beltagy, AI2

Steven Bethard, University of Arizona

Ryan Cotterell, ETH Zurich and University of Cambridge

Tanmoy Chakraborty, IIT-Delhi

Yichao Zhou, University of California, Los Angeles

**Publicity Chairs:**

Sarah Wiegrefe, Georgia Institute of Technology  
Enrico Santus, Bayer  
Peng Qi, JD AI  
Danqi Chen, Princeton

**Student Research Workshop Chairs:**

Esin Durmus, Cornell University  
Nelson Liu, Stanford University  
Vivek Gupta, University of Utah

**Student Research Workshop Faculty Advisors:**

Nanyun Peng, University of California, Los Angeles  
Yu Su, Ohio State University

**Tutorial Chairs:**

Greg Kondrak, University of Alberta  
Kalina Bontcheva, University of Sheffield  
Dan Gillick, Google Research

**Volunteer Chairs:**

Hao Cheng, Microsoft Research

**Virtual Infrastructure Committee:**

Deepak Ramachandran, Google Research

**Virtual Infrastructure Committee Advisors:**

Jan-Christoph Klie, UKP Lab, Technical University of Darmstadt  
Hao Fang, Microsoft Semantic Machines, USA  
Gisela Vallejo, Independent research

**Website Chairs:**

Ice Pasupat, Google Research  
Iulia Turc, Google Research

**Workshop Chairs:**

Bhavana Dalvi, AI2  
Mamoru Komachi, Tokyo Metropolitan University  
Michel Galley, Microsoft Research

# Senior Program Committee

## Program Chairs:

Anna Rumshisky, University of Massachusetts Lowell  
Luke Zettlemoyer, University of Washington and Facebook AI Research  
Dilek Hakkani-Tur, Amazon Alexa AI

## Senior Area Chairs and Area Chairs:

### Computational Social Science and Cultural Analytics:

SACs: David Bamman, Chenhao Tan  
ACs: Dallas Card, Paramveer Dhillon, Lucie Flek, Kenny Joseph, David Mimno, Dong Nguyen, Brendan O'Connor, Daniel Preotiuc-Pietro, Sara Tonelli, Svetlana Volkova, Justine Zhang

### Dialogue and Interactive Systems:

SACs: Maxine Eskenazi, Kallirroi Georgila, Jiwei Li  
ACs: Ryuichiro Higashinaka, Casey Kennington, Kazunori Komatani, Sungjin Lee, Pierre Lison, Jose David Lopes, Ramesh Manuvinakurike, Michael McTear, Alexandros Papangelis, Rebecca Passonneau, Gabriel Skantze, David Traum, Stefan Ultes, Nigel Ward, Jason Williams, Koichiro Yoshino, Kai Yu, Tiancheng Zhao

### Discourse and Pragmatics:

SACs: Giuseppe Carenini, Junyi (Jessy) Li  
ACs: Malihe Alikhani, Shafiq Joty, Sujian Li, Annie Louis, Deyi Xiong

### Ethics, Bias, and Fairness:

SACs: Kai-Wei Chang  
ACs: Vinodkumar Prabhakaran, Rachel Rudinger

### Green NLP:

SACs: Roy Schwartz  
ACs: Niranjan Balasubramanian, Emma Strubell

### Language Generation:

SACs: Asli Celikyilmaz, Nanyun Peng, Wei Xu  
ACs: Antoine Bosselut, Greg Durrett, Claire Gardent, Iryna Gurevych, Srinivasan Iyer, Mohit Iyyer, Yangfeng Ji, Mirella Lapata, Mike Lewis, Ankur Parikh, Laura Perez-Beltrachini, Sudha Rao, Hannah Rashkin, Mark Riedl, John Wieting

### Information Extraction:

SACs: Heng Ji, Zornitsa Kozareva, William Wang  
ACs: Gerard de Melo, Luciano Del Corro, Bhuwan Dhingra, Ruihong Huang, Lifu Huang, Arzoo Katiyar, Ni Lao, Yi Luan, Ndapa Nakashole, Vincent Ng, Thien Huu Nguyen, Qiang Ning, Sujith Ravi, Siliang Tang, Derrry Tanti Wijaya, Bishan Yang, Dian Yu, Sheng Zhang

**Information Retrieval and Text Mining:**

SACs: Eugene Agichtein, Alessandro Moschitti  
ACs: Qingyao Ai, Ahmed Hassan Awadallah, Simone Filice, Karthik Raman, Oleg Rokhlenko, Luca Soldaini, Yi Tay, Olga Uryupina, Thuy Vu, Hamed Zamani

**Interpretability and Analysis of Models for NLP:**

SACs: Yonatan Belinkov, Grzegorz Chrupala, Anders Søgaard  
ACs: Jasmijn Bastings, Allyson Ettinger, Katja Filippova, Antske Fokkens, Dieuwke Hupkes, Hung-yi Lee, Zachary Lipton, Lilja Øvrelid, Anna Rogers, Benjamin Roth, Hassan Sajjad, Wojciech Samek, Ian Tenney, Xin (Eric) Wang, Willem Zuidema

**Language Grounding to Vision, Robotics and Beyond:**

SACs: Jason Baldridge, Mohit Bansal  
ACs: Aishwarya Agrawal, Joyce Chai, Spandana Gella, David Harwath, Lisa Anne Hendricks, Angeliki Lazaridou, Karthik Narasimhan, Jesse Thomason, Mark Yatskar

**Language Resources and Evaluation:**

SACs: Sam Bowman, Matt Gardner  
ACs: Pradeep Dasigi, Kyle Gorman, Joakim Nivre, Saku Sugawara, Alane Suhr, Reut Tsarfaty, Aaron White, Adina Williams

**Linguistic Theories, Cognitive Modeling and Psycholinguistics:**

SACs: William Schuler  
ACs: Richard Futrell, Kasia Hitzenko, Emily Prud'hommeaux, Marten van Schijndel

**Machine Learning for NLP: Classification and Structured Prediction Models:**

SACs: Alan Ritter, Sasha Rush, Vivek Srikumar  
ACs: Yoav Artzi, Isabelle Augenstein, Mathieu Blondel, Shay Cohen, Caio Corro, Dan Goldwasser, Yoon Kim, Lingpeng Kong, Paul Michel, Lili Mou, Vlad Niculae, Gabriel Stanovsky, Karl Stratos, Swabha Swayamdipta

**Machine Learning for NLP: Language Modeling and Sequence to Sequence Models:**

SACs: Kevin Gimpel, Omer Levy  
ACs: Roei Aharoni, Michael Auli, Taylor Berg-Kirkpatrick, Li Dong, Tatsunori Hashimoto, Andre Martins, Matthew Peters, Colin Raffel, Rico Sennrich, Lei Yu

**Machine Translation:**

SACs: George Foster, Maja Popović, Matt Post  
ACs: Antonis Anastasopoulos, Eleftherios Avramidis, Michael Denkowski, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Paco Guzmán, Christian Hardmeier, Carolina Scarton, Dimitar Shterionov, Michel Simard, Marco Turchi, Eva Vanmassenhove, David Vilar, Rui Wang

**Multilinguality:**

SACs: Manaal Faruqui, Radu Florian  
ACs: Marine Carpuat, Orhan Firat, Dan Garrette, Katharina Kann, Sara Rosenthal, Sakriani Sakti, Ves Stoyanov, Shyam Upadhyay, Zhiguo Wang, Imed Zitouni

**NLP Applications:**

SACs: Yun-Nung Vivian Chen, Jing Huang, Nitin Madnani, Byron Wallace  
ACs: Andrew O. Arnold, Aoife Cahill, Chris Donahue, Vivek Kulkarni, Rohit Kumar, Cheng-Te Li, Yang Liu, Tristan Naumann, Hoifung Poon, Peng Qi, Nazneen Fatema Rajani, Vikram Ramanarayanan, Kevin Small, Gokhan Tur, Sowmya Vajjala, Karin Verspoor, Guangtao Wang, Youzheng Wu, Bing Xiang, Pengtao Xie, Helen Yannakoudakis, Torsten Zesch, Yuhao Zhang

**Phonology, Morphology and Word Segmentation:**

SACs: Ryan Cotterell  
ACs: Xuanjing Huang, Francis Tyers, Ekaterina Vylomova

**Question Answering:**

SACs: Danqi Chen, Hannaneh Hajishirzi  
ACs: Jonathan Berant, Jordan Boyd-Graber, Kyunghyun Cho, Eunsol Choi, Nan Duan, Kelvin Guu, Robin Jia, Daniel Khashabi, Tushar Khot, Jimmy Lin, Xiang Ren, Minjoon Seo, Pontus Stenetorp, Huan Sun

**Semantics: Lexical Semantics:**

SACs: Ivan Vulić  
ACs: Goran Glavaš, Taher Pilehvar, Roi Reichart, Steven Schockaert, Ekaterina Shutova, Aline Villavicencio

**Semantics: Sentence-level Semantics and Textual Inference:**

SACs: Daniel Cer, Douwe Kiela, Wenpeng Yin  
ACs: Heike Adel, Roy Bar-Haim, Iz Beltagy, Eduardo Blanco, Shiyu Chang, Daniel Hershcovich, Adam Poliak, Christopher G. Potts, Siva Reddy, Marek Rei, German Rigau, Cicero Nogueira dos Santos, Yangqiu Song, Yinfei Yang, Mo Yu

**Sentiment Analysis and Stylistic Analysis:**

SACs: Preslav Nakov, Diyi Yang  
ACs: Tanmoy Chakraborty, Lun-Wei Ku, Soujanya Poria, Paolo Rosso, Tamar Solorio, Aixin Sun, Arkaitz Zubiaga, Yulan He, Roman Klinger

**Speech:**

SAC: Eric Fosler-Lussier  
ACs: Erica Cooper, Yao Qian, Yuxuan Wang

**Summarization:**

SACs: Michael Elhadad, Yue Zhang  
ACs: Arman Cohan, Sebastian Gehrmann, Logan Lebanoff, Fei Liu, Pengfei Liu, Shashi Narayan, Manabu Okumura, Xiaojun Wan, Rui Zhang

**Syntax: Tagging, Chunking, and Parsing:**

SAC: Kenji Sagae  
ACs: Marco Kuhlmann, Miryam de Lhoneux, Jennifer Foster, David Chiang, Weiwei Sun, Miguel Ballesteros

**Special Theme: New Challenges in NLP:**

SAC: Kenneth Church

ACs: Chris Callison-Burch, Valia Kordoni, Sabine Schulte im Walde, Costanza Navarretta

**Conflict of Interest (COI):**

SAC: Graham Neubig

ACs: Monoji Choudhury, Anna Korhonen, Wei Lu, Kathy McKeown, Noah A. Smith

**Ethics Committee:**

Alain Couillault, Alvin Grissom II, Amanda Stent, Annemarie Friedrich, Aurélie Névéal, Paris Saclay, Chen Li, Claudia Borg, Dan Jurafsky, Darja Fišer, Dirk Hovy, George K. Acquaah-Mensah, Gilles ADDA, Jade Abbott, Jean-Yves Antoine, Jin-Dong Kim, Jingbo Xia, Jochen L. Leidner, José Ochoa-Luna, Karl Pichotta, Kevin Bretonnel Cohen, Laura Alonso Alemany, Luciana Benotti, Luis Chiruzzo, Malvina Nissim, Manny Rayner, Margot Mieskes, Maxime Amblard, Min-Yen Kan, Mona Diab, Natalie Schluter, Nyalleng Moorosi, Qingcai Chen, Robert Monarch, Ryan Georgi, Te Rutherford, Thiago Castro Ferreira, Vukosi Marivate, Wondwossen Mulugeta

**Primary Reviewers:**

Aaron Mueller, Abdalghani Abujabal, Abe Ittycheriah, Abhijit Mishra, Abhik Jana, Abhinav Gupta, Abhishek Srivastava, Abhyuday Jagannatha, Abram Handler, Ada Wan, Adam Faulkner, Adam Pauls, Adam Tsakalidis, Adhiguna Kuncoro, Aditi Chaudhary, Aditya Joshi, Aditya Kalyanpur, Aditya Mogadala, Aditya Siddhant, Adrian Benton, Adrian Iftene, Adriana Kovashka, Afshin Rahimi, Agata Savary, Agnieszka Mykowiecka, Ahmad Beirami, Ahmed Ali, Ahmed El-Kishky, Ai Ti Aw, Aida Mostafazadeh Davani, Aijun An, Aili Shen, Akash Bharadwaj, Akihiro Tamura, Akiko Aizawa, Akiko Eriguchi, Akinori Fujino, Akshay Chaturvedi, Alain Couillault, Alakananda Vempala, Alan Akbik, Alberto Lavelli, Alberto Testoni, Alborz Geramifard, Ale Tamchyna, Aleksandar Savkov, Aleksandr Drozd, Aleksandra Gabryszak, Alessandra Cervone, Alessandro Lenci, Alessandro Mazzei, Alessandro Raganato, Alessandro Sordoni, Alex Wang, Alex Warstadt, Alexander Erdmann, Alexander Fabbri, Alexander Gelbukh, Alexander Loeser, Alexander O'Connor, Alexandra Balahur, Alexandra Chronopoulou, Alexandra Lavrentovich, Alexandre Allauzen, Alexandre Denis, Alexandre Klementiev, Alexandre Rademaker, Alexey Romanov, Alexey Sorokin, Ali Ahmadvand, Alice Oh, Alina Karakanta, Alina Maria Cristea, Alisa Zhila, Alistair Willis, Alla Rozovskaya, Alon Jacovi, Alvin Grissom II, Aman Madaan, Amanda Stent, Ambedkar Dukkipati, Ameeta Agrawal, Amil Merchant, Amine Trabelsi, Aminul Islam, Amir Bakarov, Amir Feder, Amirreza Shirani, Amit Goyal, Amita Misra, Amy Siu, Anastasia Shimorina, Anastassia Kornilova, Anca Dumitrache, Andon Tchechmedjiev, Andrea Galassi, Andrea Horbach, Andrea Madotto, Andreas van Cranenburgh, Andrei Barbu, Andrew Caines, Andrew Gargett, Andrew Rosenberg, Andrey Kutuzov, Anette Frank, Angel Chang, Anietie Andy, Anil Ramakrishna, Animesh Mukherjee, Anjalie Field, Ann Bies, Ann devitt, Anna Currey, Anne Cocos, Anne Vilnat, Anne-Kathrin Schumann, Annemarie Friedrich, Annette Rios, Annika Marie Schoene, Anoop Kunchukuttan, Anoop Sarkar, Anthony Rios, Antoine Raux, Antonio Valerio Miceli Barone, Anuj Goyal, Anusha Balakrishnan, Aoife Cahill, Aparajita Haldar, Ari Rappoport, Arianna Yuan, Arijit Ghosh Chowdhury, Arjun Akula, Aron Culotta, Artur Kulmizev, Arturo Argueta, Arul Menezes, Asad Sayeed, Ashequl Qadir, Ashish Anand, Ashish Sabharwal, Ashutosh Kumar, Ashutosh Modi, Atsushi Fujita, Attapol Rutherford, Aurko Roy, Aurlie Nvol, Avanti Bhandarkar, Avia Efrat, Avik Ray, Avinesh P.V.S, Avirup Saha, Avishek Anand, Avneesh Saluja, Awais Athar, Ayah Zirikly, Aykut Erdem, Ayush Maheshwari,

Badri Narayana Patro, Baigong Zheng, Bailin Wang, Baishakhi Ray, Baiyun Cui, Balaji Ganesan, Balaji Vasan Srinivasan, Bang Liu, Baoli LI, Baolin Peng, Baosong Yang, Baotian Hu, Baoxun

Wang, Behnam Hedayatnia, Bei Chen, Bei Shi, Bei Yu, Beliz Gunel, Ben Bogin, Ben Hachey, Ben Hutchinson, Ben Peters, Ben Zhou, Benjamin Heinzerling, Benjamin Marie, Bernhard Kratzwald, Bertrand Higgy, Besim Kabashi, Bhargavi Paranjape, Bhuvana Ramabhadran, Biao Zhang, Bill Byrne, Bill Yuchen Lin, Billy Chiu, Bin Liang, Bin Wang, Bing Bai, Bing Zhao, Binod Gyawali, Bo Wang, Bo-Hsiang Tseng, Bogdan Babych, Boliang Zhang, Bonaventura Coppola, Bonnie Webber, Bowei Zou, Bowen Wu, Bowen Zhang, Boxin Wang, Boyuan Pan, Brian Davis, Brian Larson, Brian Riordan, Bridget McInnes, Bruno Martins, Bryan Routledge, Burcu Can, Byeongchang Kim, Byron Galbraith,

C. Anton Rytting, Caiwen Ding, Camilo Thorne, Can Xu, Canasai Kruengkrai, Caren Han, Carina Silberer, Carlos Escolano, Carlos Ramisch, Carolin Lawrence, Carolyn Ros, Catherine Finegan-Dollak, Catherine Havasi, Cecilia Ovesdotter Alm, Cem Rifki Aydin, Chaitanya Malaviya, Chaitanya Shivade, Chandrahas ., Chang Li, Chang Xu, Changxing Wu, Changzhi Sun, Chao Wang, Chao Zhao, Chao-Chun Hsu, Chao-Chun Liang, Charles Dognin, Charles Jochim, Charles Lovering, Chaya Liebeskind, Chen Liang, Chen Liu, Chen Zhang, Chen Zhao, Chen-Tse Tsai, Cheng Li, Chenghao Liu, Chenghao Yang, Chengyu Wang, Chengzhi Zhang, Chenhua Chen, Chenhui Chu, Chenliang Li, Chenwei Zhang, Chenyang Huang, Chester Holtz, Chia-Ying Lee, Chieh-Yang Huang, Chien-Sheng Wu, Chikara Hashimoto, Chikashi Nobata, Chinnadhurai Sankar, Chiori Hori, Chiraag Lala, Chitta Baral, Chiyu Zhang, Chong Min Lee, Chongyang Tao, Chris Alberti, Chris Brew, Chris Brockett, Chris Develder, Chris Hokamp, Chris Kedzie, Chris Quirk, Christian M. Meyer, Christian Wurm, Christo Kirov, Christoph Tillmann, Christophe Cerisara, Christophe Gravier, Christopher Homan, Christos Christodoulopoulos, Chuan-Ju Wang, Chuang Fan, Chuanyi Li, Chuhan Wu, Chung-Chi Chen, Chung-Chi Huang, Chunting Zhou, Chunyang Xiao, Chunyuan Li, Claire Bonial, Claire Gardent, Clara Vania, Clare Voss, Claudia Borg, Claudia Schulz, Claudio Greco, Colin Cherry, Colin Lockard, Constantin Orasan, Cornelia Caragea, Cory Shain, Courtney Napoles, Cuiyun Gao, Cuong Hoang, Cyril Goutte,

Dai Quoc Nguyen, Daichi Mochihashi, Daisuke Kawahara, Damien Nouvel, Dan Goldwasser, Dan Jurafsky, Dan Simonson, Dan Su, Daniel Dahlmeier, Daniel Dakota, Daniel Deutsch, Daniel Fried, Daniel Hardt, Daniel King, Daniel Stein, Daniela Moctezuma, Daniele Bonadiman, Daniil Sorokin, Danish Contractor, Danish Pruthi, Danni Ma, Danqing Wang, Danushka Bollegala, Daojian Zeng, Daoyuan Chen, Dario Bertero, Dario Stojanovski, Darja Fier, Darsh Shah, Dat Quoc Nguyen, David Adelani, David Jurgens, David Kauchak, David M. Howcroft, David Martins de Matos, David McClosky, David Schlangen, David Stap, David Talbot, David Vandyke, David Vilares, Davide Picca, Davis Liang, Daya Guo, Dayiheng Liu, Debanjan Mahata, Debela Gemechu, Debjit Paul, Debora Nozza, Deepak P, Dejing Dou, Delphine Bernhard, Deng Cai, Denis Filimonov, Denis Newman-Griffis, Deniz Zeyrek, Derek F. Wong, Devamanyu Hazarika, Devendra Sachan, Dezhi Hong, Dhananjay Ram, Dheeraj Rajagopal, Di Jin, Di Lu, Di Wang, Diana Inkpen, Diana Santos, Diane Napolitano, Dianqi Li, Diego Frassinelli, Diego Molla, Diego Moussallem, Dinesh Raghu, Dingcheng Li, Dinghan Shen, Dingquan Wang, Dipanjan Das, Dipankar Das, Diptesh Kanojia, Dirk Hovy, Divyansh Kaushik, Diya Li, Dmitry Ustalov, Dominik Schlechtweg, Don Tuggener, Dong Yu, Dong Zhou, Donghong Ji, Dongkuan Xu, Dongsheng Wang, Dongyan Zhao, Dongyu Zhang, Donna Byron, Dushyant Singh Chauhan, Duygu Ataman,

Edison Marrese-Taylor, Edoardo Maria Ponti, Eduard Dragut, Edwin Simpson, Efstathios Stamatatos, Egon Stemle, Ehsan Abbasnejad, Ehsan Kamaloo, Ehsan Shareghi, Ehsaneddin Asgari, Eiji Aramaki, Ekaterina Garmash, Ekaterina Kochmar, Ekaterina Lapshinova-Koltunski, Elad Segal, Elena Kochkina, Elena Voita, Eleonora Giunchiglia, Elisa Ferracane, Elisabetta Fersini, Elizabeth Clark, Eljko Agi, Ella Rabinovich, Ellie Pavlick, Elliott Ash, Emanuele Bastianelli, Emerson Paraiso, Emilia Apostolova, Emily Sheng, Emmanuele Chersoni, Enamul Hoque, Enrica Troiano, Enrique Amigo, Eraldo Fernandes, Eric Bell, Eric Nichols, Eric Wallace, Erick Fonseca, Erik Cambria, Erik Velldal, Erkut Erdem, Ernie Chang, Erwan Moreau, Ethan A. Chi, Ethan Perez, Ethan Selfridge, Eugene Kharitonov, Eva Maria Vecchi, Evgeny Matusov, Evgeny Stepanov,

Fabien Cromieres, Fabio Massimo Zanzotto, Fabio Tamburini, Fabrizio Silvestri, Fan Bai, Fandong Meng, Fang Kong, Fangtao Li, Farah Benamara, Farah Nadeem, Farhana Ferdousi Liza, Farig Sadeque, Fatiha Sadat, Fei Cheng, Fei Li, Fei Sun, Feiliang Ren, Felix Sasaki, Felix Stahlberg, Feng Ji, Feng Nie, Ferhan Ture, Fernando Alva-Manchego, Florian Boudin, Florian Kunneman, Florian Metze, Frances Yung, Francesco Barbieri, Francesco Corcoglioniti, Francis Ferraro, Francisco Casacuberta, Fred Popowich, Frederic Bechet, Fumiyo Fukumoto, Fuxiang Chen,

Gabriel Gordon-Hall, Gabriel ILLOUZ, Gabriele Pergola, Gaku Morio, Ganbin Zhou, Ganesh Jawahar, Ganesh Ramakrishnan, Garrett Nicolai, Gayatri Bhat, Geeticka Chauhan, Genta Indra Winata, Georg Rehm, George Acquaaah-Mensah, Gerald Penn, Gerasimos Lampouras, Gerhard Heyer, Gianluca Lebani, Gianni Barlacchi, Giannis Bekoulis, Giannis Nikolentzos, Gil Rocha, Gilles Adda, Gina-Anne Levow, Giorgio Maria Di Nunzio, Giovanni Campagna, Girish Palshikar, Giulia Venturi, Giuseppe Attardi, Giuseppe Castellucci, Gongbo Tang, Gregor Leusch, Grzegorz Kondrak, Guan-Lin Chao, Guangneng Hu, Guangyou Zhou, Guanyi Chen, Guenter Neumann, Guillaume Jacquet, Guillaume Wisniewski, Guillermo Garrido, Gunhee Kim, Guodong Long, Guohong Fu, Guokun Lai, Guoshun Nan, Guoyin Wang, Gurpreet Lehal, Gustavo Aguilar, Gustavo Henrique Paetzold, Guy Feigenblat, Guy Rosin, Gyuwon Kim,

H. Andrew Schwartz, Hai Leong Chieu, Hai Wang, Haibo Ding, Haibo Li, Haichao Zhu, Hainan Xu, Haitao Mi, Haithem Afli, Haizhou Li, Halil Kilicoglu, Hamdy Mubarak, Hamman Samuel, Hanan Aldarmaki, Hanjie Chen, Hanna Bechara, Hanna Suominen, Hannes Schulz, Hanqi Jin, Hanumant Redkar, Hanwen Zha, Hao Cheng, Hao Fang, Hao Peng, Hao Tan, Hao Tang, Hao Yuan, Hao Zhu, Haohan Wang, Haoming Jiang, Haoran Li, Haoran Zhang, Haoruo Peng, Haoxi Zhong, Haoyang Wen, Haoyu Wang, Haoyue Shi, Hardy Hardy, Harry Bunt, Harsh Jhamtani, Harsh Trivedi, Hatem Haddad, Haw-Shiuan Chang, Hayato Kobayashi, Heba Elfardy, Helena Caseli, Helena Moniz, Hen-Hsen Huang, Hend Al-Khalifa, Heng Wang, Hengyi Cai, Henning Wachsmuth, Henry Elder, Herve Dejean, Hieu Hoang, Hila Gonen, Hiroaki Sugiyama, Hiroki Ouchi, Hiroshi Echizen-ya, Hiroshi Noji, Hiroya Takamura, Hisami Suzuki, Hong Wang, Hongliang Fei, Hongming Zhang, Hongshen Chen, Hongxia Jin, Hongyu Gong, Hongzhao Huang, Hou Pong Chan, Hsin-Hsi Chen, Hsin-Min Wang, Huaishao Luo, Huajie Shao, Huichao Xue, Huijun Zhang, Huy Nguyen, Hyeju Jang, Hyung Won Chung,

I-Ta Lee, Iacer Calixto, Ian Stewart, Ignacio Iacobacci, Igor Shalyminov, Ikuya Yamada, Ilias Chalkidis, Ines Rehbein, Ingmar Weber, Ioannis Konstas, Ion Androutopoulos, Irene Russo, Iria da Cunha, Irina Temnikova, Irwin King, Iryna Gurevych, Isaac Caswell, Isar Nejadgholi, Ismail El Maarouf, Ivan Habernal, Ivan Vladimir Meza Ruiz, Ivan Yamshchikov, Izzeddin Gur,

Jaap Jumelet, Jack Hessel, Jacob Andreas, Jacob Eisenstein, Jad Kabbara, Jade Abbott, Jaegul Choo, Jaemin Cho, Jaime Lorenzo-Trueba, Jamal Abdul Nasir, James Fan, James H. Martin, James Thorne, Jamin Shin, Jan Alexandersson, Jan Botha, Jan-Christoph Klie, Jason Wei, Javid Ebrahimi, Jayant Krishnamurthy, Jean Maillard, Jean-Yves Antoine, Jeessoo Bang, Jeff Mitchell, Jekaterina Novikova, Jelena Mitrovi, Jelke Bloem, Jen-Tzung Chien, Jena D. Hwang, Jennifer D'Souza, Jennifer Williams, Jeremy Barnes, Jesse Mu, Jesse Thomason, Jessica Ouyang, Jey Han Lau, Ji He, Ji Xin, Ji-Ung Lee, Jia-Chen Gu, Jiaao Chen, Jiachang Liu, Jiacheng Xu, Jiaji Huang, Jiali Zeng, Jiaming Shen, Jiaming Xu, Jian Guan, Jian Li, Jian Ni, Jian Sun, Jian-Yun Nie, Jianfei Yu, Jiangming Liu, Jiangtong Li, Jianhua Yuan, Jianmo Ni, Jianpeng Cheng, Jianqiang Ma, Jiaqi Guo, Jiasen Lu, Jiawei Zhou, Jichuan Zeng, Jie Cao, Jie Fu, Jie Hao, Jie Yang, Jie Zhou, Jiebo Luo, Jieyu Zhao, Jill Burstein, Jimmy Xiangji Huang, Jin Wang, Jin-Dong Kim, Jin-Ge Yao, Jin-Woo Chung, JinYeong Bak, Jindich Helcl, Jindich Libovick, Jing Liu, Jing Lu, Jing Wang, Jingang Wang, Jingbo Shang, Jingbo Xia, Jingfeng Yang, Jingjing Li, Jingzhou Liu, Jinho D. Choi, Jinhua Du, Jinhyuk Lee, Jirka Hana, Jiwei Tan, Joachim Wagner, Jochen L. Leidner, Joel Tetreault, Joern Wuebker, Johannes Bjerva, Johannes Hoffart, Johannes Leveling, John Chen, John Conroy, John Hewitt, John Lawrence, John Lee, John Morris, John P. Lalor, John P. McCrae, John Pavlopoulos, Jon Chamberlain, Jonas Pfeiffer, Jonathan Clark, Jonathan Dunn, Jonathan Herzig, Jonathan K.



Kummerfeld, Jonathan May, Jonathan Schler, Jong Park, Jong-Hoon Oh, Jonggu Kim, Joo-Kyung Kim, Joonsuk Park, Jose Ochoa-Luna, Josef Ruppenhofer, Josef van Genabith, Josep Crego, Joseph Le Roux, Joshua Maynez, Juan Pino, Juan Soler, Judita Preiss, Julia Hockenmaier, Julia Ive, Julia Kreutzer, Julian Brooke, Julian Michael, Julien Kloetzer, Julien PEREZ, Julien Tourille, Jun Araki, Jun Chen, Jun Suzuki, Jun Wang, Jun Xie, Jun Xu, Jun Yan, Jun Zhu, Jun-Ping Ng, Jung-Woo Ha, Junhui Li, Junjie Hu, Junlan Feng, Junnan Zhu, Junsheng Zhou, Junta Mizuno, Juntao Li, Juntao Yu, Junwei Bao, Junwen Duan, Junxian He, Juri Opitz, Jyun-Yu Jiang,

Kai Chen, Kai Hong, Kai Hui, Kai Sun, Kai Zhao, Kaiqiang Song, Kalpa Gunaratna, Kalpesh Krishna, Kam-Fai Wong, Kamal Sarkar, Karan Aggarwal, Kareem Darwish, Karl Pichotta, Karl Stratos, Karthik Gopalakrishnan, Kartik Goyal, Katerina Margatina, Katherine Keith, Katherine Stasaski, Katrin Erk, Katrin Kirchhoff, Katsuhiko Hayashi, Katsuhito Sudoh, Kaylee Burns, Kazunari Sugiyama, Kazutoshi Shinoda, Keh-Yih Su, Kehai Chen, Keisuke Sakaguchi, Keith VanderLinden, Kellie Webster, Kemal Ofazer, Ken Barker, Kenichi Iwatsuki, Kenji Imamura, Kenneth Heafield, Kenny Zhu, Kentaro Inui, Kentaro Torisawa, Kenton Lee, Kenton Murray, Kerstin Jung, Kevin Cohen, Kevin Duh, Kevin Lin, Kevin Seppi, Kewei Tu, Kiem-Hieu Nguyen, Kilian Evang, Kishaloy Halder, Kiyonori Ohtake, Kiyotaka Uchimoto, Koki Washio, Kokil Jaidka, Kordula De Kuthy, Kotaro Funakoshi, Koustuv Sinha, Krishnapriya Vishnubhotla, Kriste Krstovski, Kristen Johnson, Kristjan Arumae, Kuan-Yu Chen, Kugatsu Sadamitsu, Kumiko Tanaka-Ishii, Kun Qian, Kun Xu, Kunpeng Zhang, Kushal Kafle, Kuzman Ganchev, Kwan Hui Lim, Kyle Johnson, Kyle Lo, Kyle Richardson, Kyu Han, Kyungjae Lee,

Lambert Mathias, Lan Du, Lanbo She, Larry Moss, Lasha Abzianidze, Laura Alonso Alemany, Laura Rimell, Laurent Besacier, Le An Ha, Le Sun, Lea Frermann, Lee Becker, Lei Cao, Lei Hou, Lei Sha, Lei Shu, Lei Zhang, Leo Wanner, Leonardo F. R. Ribeiro, Leshem Choshen, Li Dong, Liane Guillou, Liang Wang, Liang-Chih Yu, Liangyou Li, Lianhui Qin, Libin Liu, Libo Qin, Lidia Pivovarova, Lidong Bing, Lieke Gelderloos, Lifeng Jin, Lifu Tu, Lijun Lyu, Lijun Wu, Likun Qiu, Liling Tan, Lin Chen, Lin Gui, Lin Zhao, Lin Zheng, Liner Yang, Ling Liu, Lingfei Wu, Lingjia Deng, Lingyu Gao, Linyang Li, Lis Pereira, Lisa Bauer, Lisheng Fu, Liye Fu, Lizhen Qu, Longyue Wang, Lubomir Ivanov, Luca Benedetto, Lucia Donatelli, Luciana Benotti, Lucien Carroll, Lucy Vanderwende, Luigi Di Caro, Luis Chiruzzo, Luis Espinosa Anke, Luis Gerardo Mojica de la Vega, Lun-Wei Ku, Lung-Hao Lee, Luwen (Vivian) Huangfu,

M. Amin Farajian, MIN XIAO, Maarten Sap, Maciej Ogrodniczuk, Maciej Piasecki, Magnus Sahlgren, Mahesh Joshi, Maja Popovi, Majid Laali, Majid Yazdani, Makoto Miwa, Maksim Podkorytov, Malihe Alikhani, Malvina Nissim, Mamoru Komachi, Manabu Torii, Manasi Patwardhan, Mandar Joshi, Manex Agirrezabal, Manik Bhandari, Manny Rayner, Mantong Zhou, Manuel Montes, Manuel R. Ciosici, Manuela Speranza, Maor Ivgi, Marc Tomlinson, Marcello Federico, Marco Antonio Sobrevilla Cabezudo, Marco Damonte, Marco Del Tredici, Marco Di Giovanni, Marco Passarotti, Marcos Garcia, Marcos Goncalves, Marcos Treviso, Marcos Zampieri, Mareike Hartmann, Margot Mieskes, Maria Antoniak, Maria Leonor Pacheco, Maria Liakata, Maria Nadejde, Maria Pia di Buono, Mariana Neves, Marianna Martindale, Mariano Felice, Marie Candito, Marie-Jean Meurs, Marina Fomicheva, Marina Litvak, Marina Santini, Mario Graff, Marius Mosbach, Marius Pasca, Marjan Hosseinia, Mark Cieliebak, Mark Granroth-Wilding, Mark Hasegawa-Johnson, Mark Hopkins, Mark Last, Mark Steedman, Mark-Jan Nederhof, Markus Egg, Marta Tatu, Marti A. Hearst, Martin Gleize, Maryam Fazel-Zarandi, Masaaki Nagata, Masaaki Nishino, Masaharu Yoshioka, Masahiro Kaneko, Masao Utiyama, Masashi Shimbo, Masashi Yoshikawa, Masato Mita, Masoud Rouhizadeh, Massimo Nicosia, Massimo Poesio, Masud Mosh-taghi, Mathieu Lafourcade, Matteo Negri, Matthew McDermott, Matthew Mulholland, Matthew Purver, Matthew Sims, Matthias Huck, Matthieu Labeau, Maud Ehrmann, Mauro Cettolo, Max Bartolo, Max Silberztein, Maxime Amblard, Maxime Peyrard, Maximin Coavoux, Mayank Singh, Md Arafat Sultan, Md Faisal Mahbub Chowdhury, Mehdi Rezagholizadeh, Mehrad Moradshahi, Mehrnoosh Sadrzadeh, MeiHua Chen, Meishan Zhang, Melissa Roemmele, Melvin Johnson, Meng Fang, Meng Zhou, Mengjie Zhao, Mengnan Du, Mengting Hu, Mengwen Liu, Mengzhou Xia,

Menno van Zaanen, Meriem Beloucif, Mert Kilickaya, Messina Enza, Micha Elsner, Michael A. Laurenzano, Michael Bloodgood, Michael Bugert, Michael Collins, Michael Elhadad, Michael Flor, Michael Gamon, Michael Mohler, Michael Neumann, Michael Regan, Michael Roth, Michael Sejr Schlichtkrull, Michael Wayne Goodman, Michael Wiegand, Michael Zock, Michal Ptaszynski, Miguel A. Alonso, Miguel Domingo, Mihael Arcan, Mihir Kale, Miikka Silfverberg, Mika Juuti, Mikhail Khodak, Mikko Kurimo, Min-Yen Kan, Ming Liu, Mingbo Ma, Mingda Chen, Mingxuan Wang, Mingyue Shang, Minwoo Jeong, Minwoo Lee, Miriam Butt, Miriam R L Petruck, Mladen Karan, Mohamad Hardyman Barawi, Mohamed Morchid, Mohammad Akbari, Mohammad Javad Hosseini, Mohammad Salameh, Mohsen Mesgar, Momchil Hardalov, Mona Diab, Montse Cuadros, Mor Geva, Mostafa Abdou, Mounica Maddela, Mourad Abbas, Mozhi Zhang, Mrcis Pinnis, Mrinmaya Sachan, Muhammad Abdul-Mageed, Muhammad Humayoun, Muhao Chen, Muhua Zhu, Munindar Singh,

Nabiha Asghar, Nader Akoury, Nadi Tomeh, Nadir Durrani, Nancy Chen, Nancy Ide, Naoaki Okazaki, Naoki Otani, Naoki Yoshinaga, Naoya Inoue, Nasredine Semmar, Nasser Zalmout, Natalia Grabar, Natalia Loukachevitch, Natalia Vanetik, Natalie Parde, Natalie Schluter, Nate Blaylock, Naveen Arivazhagan, Navid Rekabsaz, Navin Goyal, Nayeon Lee, Nedjma Ousidhoum, Nelson F. Liu, Nemanja Djuric, Ngoc Phuoc An Vo, Ngoc Thang Vu, Nianzu Ma, Nicholas Andrews, Nicholas FitzGerald, Nicholas Monath, Nicola De Cao, Nicole Peinelt, Nikhil Krishnaswamy, Nikita Kitaev, Nikita Moghe, Niko Partanen, Nikola I. Nikolov, Nikolaos Aletras, Nikolaos Lagos, Nikolaos Pappas, Nikos Papasarantopoulos, Niloofar Safi Samghabadi, Nina Poerner, Nina Tahmasebi, Ningyu Zhang, Nitendra Rajput, Nithum Thain, Nitin Ramrakhiyani, Nitish Gupta, Noah Weber, Noyal B. Niraula, Nona Naderi, Nora Hollenstein, Noriki Nishida, Noriyuki Kojima, Nut Limsopatham, Nyalleng Moorosi,

Oana Cocarascu, Oana Inel, Ofir Press, Oier Lopez de Lacalle, Oladimeji Farri, Oleg Vasilyev, Ondej Duek, Oren Barkan, Ori Ram, Ori Shapira, Oshin Agarwal,

Pablo Duboue, Pablo Loyola, Panupong Pasupat, Paolo Torroni, Paramita Mirza, Pararth Shah, Paria Jamshid Lou, Parminder Bhatia, Parnia Bahar, Parul Awasthy, Pascal Denis, Pasquale Minervini, Pat Verga, Patrick Ehlen, Patrick Lewis, Patrick Paroubek, Patrick Saint-Dizier, Paul Buitelaar, Paul Cook, Paul Pu Liang, Pavan Kapanipathi, Pavankumar Reddy Muddireddy, Pavel Pecina, Pavel Petrushkov, Pawan Goyal, Pedro Henrique Martins, Pedro Rodriguez, Pegah Nokhiz, Peifeng Li, Peng Liu, Peng Shi, Peng Xu, Pengcheng Yin, Pengxiang Cheng, Peter Anderson, Peter Hase, Peter Makarov, Peter Shaw, Petya Osenova, Philip Cohen, Philipp Cimiano, Philipp Dufter, Philipp Koehn, Philippe Muller, Phong Le, Phu Mon Htut, Piek Vossen, Pierpaolo Basile, Pierre Nugues, Piji Li, Ping Jian, Piush Aggarwal, Piyawat Lertvittayakumjorn, Po Hu, Pouya Pezeshkpour, Prafulla Kumar Choubey, Pramod Kaushik Mudrakarta, Prasanna Parthasarathi, Prathyusha Jwalapuram, Pratyay Banerjee, Pratyush Maini, Preethi Jyothi, Preethi Raghavan, Preethi Vaidyanathan, Priyanka Agrawal, Prodromos Malakasiotis, Przemyslaw Biecek, Pushkar Mishra,

Qi Zhang, Qian Chen, Qian Liu, Qian Yang, Qianglong Chen, Qiaozi Gao, Qing Li, Qingcai Chen, Qingming Tang, Qingqing Cao, Qingrong Xia, Qingyu Yin, Qingyu Zhou, Qingyun Wang, Qinyuan Ye, Qiongfai Xu, Quan Hung Tran, Quanzhi Li, Qun Liu,

R. Thomas McCoy, Rabih Zbib, Radhika Mamidi, Radu Tudor Ionescu, Rafael E. Banchs, Raghav Gupta, Rahul Aralikkatte, Rahul Bhagat, Rahul Goel, Rahul Jha, Raj Dabre, Rajarshi Bhowmik, Rajhans Samdani, Rajkumar Pujari, Rakesh Verma, Ralf Krestel, Ramakanth Pasunuru, Ramit Sawhney, Ramon Lopez-Cozar, Ramy Eskander, Ran Levy, Raphael Rubino, Raphael Shu, Raphael Tang, Raul Puri, Ravi Shekhar, Raymond Wong, Reid Pryzant, Renjie Zheng, Richard Johansson, Rik Koncel-Kedziorski, Rik van Noord, Rishabh Bhardwaj, Rishi Bommasani, Rishiraj Saha Roy, Ritvik Shrivastava, Rivka Levitan, Riza Batista-Navarro, Rob Malouf, Rob Voigt, Rob van der Goot, Robert Berwick, Robert Monarch, Robert Schwarzenberg, Roberto Basili, Roberto Navigli, Rodolfo Delmonte, Rodrigo Agerri, Rodrigo Wilkens, Roe Aharoni, Roger Levy, Rohan Ramanath, Rohil Verma, Rohini Srihari, Rohit Babbar, Roland Kuhn, Roma Patel, Ronen Tamari, Rotem Dror,

Rowan Zellers, Roxanne El Baff, Rudolf Rosa, Rudra Murthy, Ruey-Cheng Chen, Ruhi Sarikaya, Rui Meng, Ruizhe Li, Rumen Dangovski, Runzhe Yang, Ruochoen Xu, Ruqing Zhang, Ryan Georgi, Ryan McDonald, Ryohei Sasano, Ryu Iida, Ryuichi Takanobu,

Saab Mansour, Sabine Bergler, Sabrina J. Mielke, Sachin Mehta, Sadid A. Hasan, Saeed Najafi, Saghar Hosseini, Sahisnu Mazumder, Saif Mohammad, Salam Khalifa, Salim Roukos, Salvatore Romeo, Sam Davidson, Sam Wiseman, Samar Husain, Sameen Maruf, Sameer Bansal, Samira Shaikh, Samuel Broscheit, Samuel Louvan, Sanchit Agarwal, Sanda Harabagiu, Sandeep Mathias, Sandeep Soni, Sandro Pezzelle, Sanjay Subramanian, Sanjeev Kumar Karn, Sanqiang Zhao, Santanu Pal, Sara Meftah, Sara Stymne, Sarah Moeller, Sarah Wiegrefe, Sarthak Jain, Sascha Rothe, Satoshi Sekine, Satwik Bhattamishra, Saurabh Garg, Saurav Sahay, Sayan Ghosh, Scott Piao, Sean MacAvaney, Sean Welleck, Sebastian Ebert, Sebastian Schuster, Sebastien Delecraz, Seema Nagar, Seid Muhie Yimam, Seiichi Nakagawa, Selma Tekir, Semih Yavuz, Seniz Demir, Senthil Chandramohan, Seojin Bang, Seokhwan Kim, Seraphina Goldfarb-Tarrant, Serena Jeblee, Serena Villata, Serge Sharoff, Sergey Feldman, Sergiu Nisioi, Serra Sinem Tekiroglu, Seunghak Yu, Seunghyun Yoon, Seungjae Shin, Seungwhan Moon, Sewon Min, Shachi Paul, Shahram Khadivi, Shamil Chollampatt, Shane Steinert-Threlkeld, Shankar Kumar, Shaolin Zhu, Shaonan Wang, Sharon Small, Shauli Ravfogel, Shen Gao, Sheng Li, Sheng Zha, Shi Feng, Shigehiko Schamoni, Shijie Wu, Shikhar Vashishth, Shikib Mehri, Shimi Salant, Shinji Watanabe, Shirley Anugrah Hayati, Shivashankar Subramanian, Sho Takase, Shoaib Jameel, Shomir Wilson, Shota Sasaki, Shruti Jadon, Shruti Palaskar, Shruti Rijhwani, Shuai Tang, Shuai Wang, Shuang Liu, Shuangyin Li, Shujian Huang, Shuming Shi, Shuo Ren, Shuohang Wang, Sida I. Wang, Siddhant Garg, Siddharth Karamcheti, Siddharth Patwardhan, Siddhartha Banerjee, Simon Dobnik, Simon Keizer, Simon Mille, Sittichai Jiampojarn, Sivaji Bandyopadhyay, Siyao Peng, Siyu Liao, Slav Petrov, Sonal Gupta, Songfang Huang, Sophia Ananiadou, Sorami Hisamoto, Soravit Changpinyo, Soroush Vosoughi, Soumya Ray, Soumya Sanyal, Souvik Kundu, Spencer Whitehead, Sravan Bodapati, Sravana Reddy, Staffan Larsson, Stan Szpakowicz, Stefan Larson, Stefan Schweter, Steffen Eger, Steffen Remus, Stephan Peitz, Stephen Mayhew, Stephen Pulman, Stergos Afantenos, Steve Skiena, Steve Young, Steven Bedrick, Steven Wilson, Su Lin Blodgett, Su-Youn Yoon, Subhashini Venugopalan, Subhro Roy, Sucheta Ghosh, Suchin Gururangan, Sudeep Gandhe, Sudha Bhingardive, Sudip Naskar, Sufeng Niu, Sujatha Das Gollapalli, Sujay Kumar Jauhar, Sukanta Sen, Sumit Bhatia, Sun Kim, Sunayana Sitaram, Sunghwan Mac Kim, Sungjoon Park, Sunil Kumar Sahu, Surafel Melaku Lakew, Surangika Ranathunga, Suzanna Sia, Sven Buechel, Svitlana Vakulenko, Swarnadeep Saha, Syed Sarfaraz Akhtar, Sylvain Kahane, Sylvain Pogodalla,

T. J. Hazen, T.Y.S.S Santosh, Tadashi Nomoto, Tae Yano, Taesun Moon, Tagyoung Chung, Takuma Udagawa, Tamer Alkhouli, Tanja Samardzic, Tanmoy Chakraborty, Tanya Goyal, Tao Chen, Tao Li, Tao Shen, Tao Yu, Taraka Rama, Tariq Alhindi, Taro Watanabe, Tatyana Ruzsics, Tejas Gokhale, Tejas Srinivasan, Terra Blevins, Teruhisa Misu, Thanh-Le Ha, Thiago Castro Ferreira, Thibault Sellam, Thierry Hamon, Thomas Brovelli (Meyer), Thomas Kober, Thomas Kollar, Thomas Lippincott, Tianlu Wang, Tianrui Li, Tianyi Liu, Tianyong Hao, Tianyu Gao, Tianyu Liu, Tianyu Zhao, Tim O’Gorman, Tim Van de Cruys, Timo Baumann, Timo Schick, Timothy Baldwin, Timothy Miller, Ting Liu, Tingting Mu, Tirthankar Ghosal, Tommaso Fornaciari, Tommaso Pasini, Tommi Pirinen, Tomoyuki Kajiwara, Toms Bergmanis, Tong Liu, Tong Niu, Tong Wang, Tong Xiao, Tongfei Chen, Tongtao Zhang, Toshiaki Nakazawa, Tracy Holloway King, Trapit Bansal, Trevor Cohn, Tristan Miller, Truc-Vien T. Nguyen, Trung Bui, Tsu-Jui Fu, Tsutomu Hirao, Tsuyoshi Okita, Tsvetomila Mihaylova, Tu Vu, Tuhin Chakraborty,

Udo Hahn, Udo Kruschwitz, Utpal Garain,

V.G.Vinod Vydiswaran, Valentin Barriere, Valentin Malykh, Valeria de Paiva, Varun Gangal, Varun Kumar, Vasileios Lamos, Vassilina Nikoulina, Venelin Kovatchev, Vered Shwartz, Verena Lyding, Verginica Barbu Mititelu, Veronika Vincze, Veronique Hoste, Victoria Bobicev, Vidas Daudaravicius, Vidhisha Balachandran, Viet-An Nguyen, Vihan Jain, Vijayaditya Peddinti, Vikas Yadav, Vikram

Ramanarayanan, Viktor Hangya, Vinay Rao, Vincent Vandeghinste, Vinit Ravishankar, Vittorio Castelli, Vivek Gupta, Vivek Kulkarni, Vivek Kumar Rangarajan Sridhar, Viviana Patti, Vladimir Eidelman, Vukosi Marivate,

Wai Lam, Wajdi Zaghouni, Wanxiang Che, Wasi Ahmad, Wayne Xin Zhao, Wei Gao, Wei Wang, Wei Wei, Wei Wu, Wei Yang, Wei Zhang, Wei Zhao, Wei-Nan Zhang, Weifeng Su, Weihua Luo, Weiwei Cheng, Weiyan Shi, Weiyang Liu, Weiyue Wang, Wen Wang, Wen Xiao, Wen Zhang, Wencan Luo, Wenduan Xu, Wenge Rong, Wenhan Luo, Wenhan Xiong, Wenhua Chen, Wenya Guo, Wenzheng Zhang, William Boag, William Gantt, Winston Wu, Wojciech Kryscinski, Wolfgang Macherey, Wolfgang Maier, Wray Buntine, Wuwei Lan,

Xavier Carreras, Xia Cui, Xian Li, Xiang Dai, Xiang Gao, Xiang Li, Xiang Lorraine Li, Xiang Yue, Xiang Zhou, Xiangyang Zhou, Xiangyu Duan, Xianpei Han, Xiao Ling, Xiao Liu, Xiaochi Wei, Xiaochuang Han, Xiaojun Chang, Xiaojun Quan, Xiaolei Huang, Xiaotong Zhang, Xilun Chen, Xin Li, Xin Lv, Xin Zheng, Xinchu Chen, Xing Niu, Xingdi Yuan, Xingshan Zeng, Xingxing Zhang, Xingyi Song, Xintong Li, Xinyi Wang, Xinyu Dai, Xinyu Hua, Xinyuan Zhang, Xipeng Qiu, Xiujun Li, Xiuying Chen, Xu Han, Xu Tan, Xuancong Wang, Xuchao Zhang, Xuezhe Ma,

Yiping Song, Yacine Jernite, Yadollah Yaghoobzadeh, Yan Song, Yan Wang, Yan Zhang, Yanai Elazar, Yanbin Lu, Yang Deng, Yang Gao, Yang Li, Yang Zhao, Yang Zhou, Yankai Lin, Yannick Versley, Yannis Korkontzelos, Yanran Li, Yansong Feng, Yanyan Zou, Yanzeng Li, Yao Meng, Yao Zhao, Yaowei Zheng, Yash Kumar Lal, Yashar Mehdad, Yasumasa Onoe, Yawei Sun, Yazhou Zhang, Yee Seng Chan, Yen-Chen Wu, Yen-Chun Chen, Yequan Wang, Yevgen Matushevych, Yevgeniy Puzikov, Yeyun Gong, Yi Cai, Yi Yang, Yi Zhang, Yi-An Lai, Yi-Ting Huang, Yibo Sun, Yichao Lu, Yichun Yin, Yiding Hao, Yifan Gao, Yifan He, Yifan Hu, Yifan Zhang, Yilin Shen, Yimai Fang, Yiming Cui, Yin-Wen Chang, Ying Lin, Yingwei Xin, Yinhe Zheng, Yisong Miao, Yitong Li, Yixin Cao, Yixin Nie, Yizhong Wang, Yllias Chali, Yo Ehara, Yogarshi Vyas, Yonatan Bitton, Yoshihiko Hayashi, Yoshihiko Suhara, Yoshimasa Tsuruoka, Yoshinari Fujinuma, Yoshinobu Kano, Younes Samih, Youngja Park, Yu Chen, Yu Su, Yu Wan, Yu Zhang, Yuan Luo, Yuan Zhang, Yuan-Fang Li, Yuanbin Wu, Yuanliang Meng, Yuanpeng Li, Yuchen Lu, Yue Chen, Yue Dong, Yue Gu, Yuen-Hsien Tseng, Yuexin Wu, Yufang Hou, Yuhui Zhang, Yuichiroh Matsubayashi, Yuji Matsumoto, Yujia Bao, Yuki Arase, Yulia Tsvetkov, Yulong Chen, Yumo Xu, Yun-hsuan Sung, Yunfang Wu, Yung-Chun Chang, Yuning Mao, Yuntian Deng, Yury Zemlyanskiy, Yusuke Miyao, Yuval Pinter, Yuxiang Wu, Yuxuan Lai, Yuxuan Wang, Yves Lepage, Yves Scherrer, Yvonne Adesam,

Zdenk abokrtsk, Ze Yang, Zeerak Waseem, Zewei Chu, Zhanming Jie, Zhao Yan, Zhaohui Wu, Zhaojiang Lin, Zhaopeng Tu, Zhe Gan, Zhe Zhang, Zhen Hai, Zhen Wang, Zhen Xu, Zheng Li, Zhenghao Liu, Zhengli Zhao, Zhengxian Gong, Zhengyuan Liu, Zhengzhong Liu, Zhenhua Ling, Zhenxi Lin, Zhi-Hong Deng, Zhijian Liu, Zhijing Jin, Zhiliang Tian, Zhirui Zhang, Zhisong Zhang, Zhiting Hu, Zhiwei Deng, Zhiyu Chen, Zhong Su, Zhongkai Sun, Zhongqing Wang, Zhouhan Lin, Zhouxing Shi, Zhu Fengda, Zhuang Liu, Zhun Liu, Zhunchen Luo, Zhuolin Jiang, Zhuoren Jiang, Zhuosheng Zhang, Zhuoxuan Jiang, Zi Lin, Zi Yang, Zi-Yi Dou, Zichao Li, Zied Bouraoui, Zihan Liu, Zijian Wang, Ziming Huang, Zining Zhu, Ziqiang Cao, Zirui Wang, Zita Marinho, Ziyi Yang, Ziyu Yao

### Secondary Reviewers:

Abhisek Chakrabarty, Aditya Bhargava, Agathe Balayn, Akshay Chaturvedi, Aku Rouhe, Alda Mari, Amir Soleimani, Ana-Maria Bucur, Andrea Amelio Ravelli, Andrey Sakryukin, Archita Pathak, Arkil Patel, Arnob Mallik, Ashish Shenoy, Avi Caciularu,

Baikjin Jung, Ben Eyal, Bhargav Kurma, Bianca Scarlini, Bing Tian, Bishal Santra, Bradley Hauer, C.S. Bahushruth, Carlos Martinez, ChaeHun Park, Chandrashekar Konda, Chen (Patrick) Pei, Chen Zhang, Chenchen Xu, Chenglei Si, Chi-kiu Lo, Christian Angel, Chung-Wei Hang,

Dan Li, Daniel Bis, Dejan Porjazovski, Di Wu, Diego Maupomé, Dingcheng Li, Dongfang Li, Dr. Sachin Pawar,  
Eric Lehman, Eugene Jang,  
Fanny Rancourt,  
Gabriel Bretschner, Garrett Nicolai, Gaurish Thakkar, Gennaro Nolano, Gong Zhang, Guanhua Zhang,  
Hali Lindsay, Han He, Haolan Zhan, Haolin Song, Hiroshige Aoki, Hong Lei Guo, Hongchang Bao, Hongyin Tang, Hossein Rouhizadeh, Huy Vu,  
Ikumi Yamashita, Itzik Malkiel,  
James Finch, Jeremy Lei Liu, Jiayuan Mao, Jing Chen, Johannes Deleu, Joy Mahapatra, Juho Leinonen, Jungsoo Park,  
Kevin Stowe, Klim Zaporojets, Kobra Amiri,  
Leonard Dahlmann, Liang Ding, Lili Wang, Linh Pham, Linyang Li, Liyan Xu,  
Machel Reid, Matt McNeil, Matthew Johnson, Matthew Matero, Maury Courtland, Maxime D. Armstrong, Maximilian Spliethöver, Mayur Patidar, Małgorzata Marciniak, Mehdi Rezaee, Miao Li, Michael Abaho, Michele Bevilacqua, Mingfei Han, Mingjie Li, Mingli Wu, Mohamed Lichouri, Molly Gia, Mujeen Sung, Mykola Lukashchuk,  
Nan Shao, Nastaran Babanejad, Nils Feldhus, None,  
Onkar Pandit, Ori Katz, Orion Weller,  
Pere-Lluis Hugué Cabot,  
Rocco Tripodi, Roy Eisenstadt, Rudolf Schneider, Ruibo Liu,  
Sabin Kaffle, Sameer Gupta, Sanket Vaibhav Mehta, Satwik Bhattamishra, Savitha Sam Abraham, Sean S. Yi, Semere Kiros Bitew, Shanshan Huang, Sheng Cao, Shi Ke, Shutian Ma, Siyu Ren, Sougata Saha, Souvik Das, Sowmya S Sundaram,  
Tahmid Rahman, Tianshi Che, Timour Igamberdiev, Ting Hua, Twisampati Sarkar,  
Vaibhav Garg, Vitalii Zhelezniak,  
Wang Huimin, Wei-Fan Chen, Weicheng Ma, Wonsuk Yang,  
Xanh Thi Ho, Xenia Klinge, Xiangyu Yang, Xiaobo Guo, Xibin Gao, Xin Chen, Xuan Zhao,  
Yanpeng Zhao, Yaocheng Gui, Yenchang Hsu, Yicheng Wang, Yingya Li, Yingyi Zhang, Yinliang Wang, Yisi Sang, Yiyi Liu, Yizhu Liu, Yongjing Yin, Yuanhe Tian, Yue Wang, Yuntao Li, Yuqing Xie,  
Zhenzhen Li, Zhaofeng Wu, Zhengyuan Liu, Shi Ke, Zhihong Chen, Zhihua Wen, Zhijia Chen, Zhijie Zhou, Zhuofeng Wu, Zhuoyi Wang



## Table of Contents

<i>Knowledge Router: Learning Disentangled Representations for Knowledge Graphs</i> Shuai Zhang, Xi Rao, Yi Tay and Ce Zhang .....	1
<i>Distantly Supervised Relation Extraction with Sentence Reconstruction and Knowledge Base Priors</i> Fenia Christopoulou, Makoto Miwa and Sophia Ananiadou .....	11
<i>Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks</i> Minh Van Nguyen, Viet Lai and Thien Huu Nguyen .....	27
<i>Abstract Meaning Representation Guided Graph Encoding and Decoding for Joint Information Extraction</i> Zixuan Zhang and Heng Ji .....	39
<i>A Frustratingly Easy Approach for Entity and Relation Extraction</i> Zexuan Zhong and Danqi Chen .....	50
<i>Event Time Extraction and Propagation via Graph Attention Networks</i> Haoyang Wen, Yanru Qu, Heng Ji, Qiang Ning, Jiawei Han, Avi Sil, Hanghang Tong and Dan Roth	62
<i>Probing Word Translations in the Transformer and Trading Decoder for Encoder Layers</i> Hongfei Xu, Josef van Genabith, Qiuhui Liu and Deyi Xiong .....	74
<i>Mediators in Determining what Processing BERT Performs First</i> Aviv Slobodkin, Leshem Choshen and Omri Abend .....	86
<i>Automatic Generation of Contrast Sets from Scene Graphs: Probing the Compositional Consistency of GQA</i> Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz and Michael Elhadad .....	94
<i>Multilingual Language Models Predict Human Reading Behavior</i> Nora Hollenstein, Federico Pirovano, Ce Zhang, Lena Jäger and Lisa Beinborn .....	106
<i>Do Syntactic Probes Probe Syntax? Experiments with Jabberwocky Probing</i> Rowan Hall Maudslay and Ryan Cotterell .....	124
<i>A Non-Linear Structural Probe</i> Jennifer C. White, Tiago Pimentel, Naomi Saphra and Ryan Cotterell .....	132
<i>Concealed Data Poisoning Attacks on NLP Models</i> Eric Wallace, Tony Zhao, Shi Feng and Sameer Singh .....	139
<i>Backtranslation Feedback Improves User Confidence in MT, Not Quality</i> Vilém Zouhar, Michal Novák, Matúš Žilinec, Ondřej Bojar, Mateo Obregón, Robin L. Hill, Frédéric Blain, Marina Fomicheva, Lucia Specia and Lisa Yankovskaya .....	151
<i>Data Filtering using Cross-Lingual Word Embeddings</i> Christian Herold, Jan Rosendahl, Joris Vanvinckenroye and Hermann Ney .....	162

<i>Improving the Lexical Ability of Pretrained Language Models for Unsupervised Neural Machine Translation</i>	
Alexandra Chronopoulou, Dario Stojanovski and Alexander Fraser .....	173
<i>Neural Machine Translation without Embeddings</i>	
Uri Shaham and Omer Levy .....	181
<i>Counterfactual Data Augmentation for Neural Machine Translation</i>	
Qi Liu, Matt Kusner and Phil Blunsom .....	187
<i>Cultural and Geographical Influences on Image Translatability of Words across Languages</i>	
Nikzad Khani, Isidora Tourni, Mohammad Sadegh Rasooli, Chris Callison-Burch and Derry Tanti Wijaya .....	198
<i>Multilingual BERT Post-Pretraining Alignment</i>	
Lin Pan, Chung-Wei Hang, Haode Qi, Abhishek Shah, Saloni Potdar and Mo Yu .....	210
<i>A Million Tweets Are Worth a Few Points: Tuning Transformers for Customer Service Tasks</i>	
Amir Hadifar, Sofie Labat, Veronique Hoste, Chris Develder and Thomas Demeester .....	220
<i>Paragraph-level Rationale Extraction through Regularization: A case study on European Court of Human Rights Cases</i>	
Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos and Prodromos Malakasiotis .....	226
<i>Answering Product-Questions by Utilizing Questions from Other Contextually Similar Products</i>	
Ohad Rozen, David Carmel, Avihai Mejer, Vitaly Mirkis and Yftah Ziser .....	242
<i>EnSidNet: Enhanced Hybrid Siamese-Deep Network for grouping clinical trials into drug-development pathways</i>	
Lucia Pagani .....	254
<i>DATE: Detecting Anomalies in Text via Self-Supervision of Transformers</i>	
Andrei Manolache, Florin Brad and Elena Burceanu .....	267
<i>A Simple Approach for Handling Out-of-Vocabulary Identifiers in Deep Learning for Source Code</i>	
Nadezhda Chirkova and Sergey Troshin .....	278
<i>Fast and Scalable Dialogue State Tracking with Explicit Modular Decomposition</i>	
Dingmin Wang, Chenghua Lin, Qi Liu and Kam-Fai Wong .....	289
<i>Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks</i>	
Nandan Thakur, Nils Reimers, Johannes Daxenberger and Iryna Gurevych .....	296
<i>SmBoP: Semi-autoregressive Bottom-up Semantic Parsing</i>	
Ohad Rubin and Jonathan Berant .....	311
<i>SGL: Speaking the Graph Languages of Semantic Parsing via Multilingual Translation</i>	
Luigi Procoppio, Rocco Tripodi and Roberto Navigli .....	325
<i>Unifying Cross-Lingual Semantic Role Labeling with Heterogeneous Linguistic Resources</i>	
Simone Conia, Andrea Bacciu and Roberto Navigli .....	338



<i>Fool Me Twice: Entailment from Wikipedia Gamification</i>	
Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger and Jordan Boyd-Graber .....	352
<i>Meta-Learning for Domain Generalization in Semantic Parsing</i>	
Bailin Wang, Mirella Lapata and Ivan Titov .....	366
<i>Aspect-Controlled Neural Argument Generation</i>	
Benjamin Schiller, Johannes Daxenberger and Iryna Gurevych .....	380
<i>Text Generation from Discourse Representation Structures</i>	
Jiangming Liu, Shay B. Cohen and Mirella Lapata .....	397
<i>APo-VAE: Text Generation in Hyperbolic Space</i>	
Shuyang Dai, Zhe Gan, Yu Cheng, Chenyang Tao, Lawrence Carin and Jingjing Liu .....	416
<i>DART: Open-Domain Structured Data Record to Text Generation</i>	
Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher and Nazneen Fatema Rajani .....	432
<i>When Being Unseen from mBERT is just the Beginning: Handling New Languages With Multilingual Language Models</i>	
Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot and Djamé Seddah .....	448
<i>Multi-Adversarial Learning for Cross-Lingual Word Embeddings</i>	
Haozhou Wang, James Henderson and Paola Merlo .....	463
<i>Multi-view Subword Regularization</i>	
Xinyi Wang, Sebastian Ruder and Graham Neubig .....	473
<i>mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer</i>	
Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua and Colin Raffel .....	483
<i>MetaXL: Meta Representation Transformation for Low-resource Cross-lingual Learning</i>	
Mengzhou Xia, Guoqing Zheng, Subhabrata Mukherjee, Milad Shokouhi, Graham Neubig and Ahmed Hassan Awadallah .....	499
<i>Open Domain Question Answering over Tables via Dense Retrieval</i>	
Jonathan Herzig, Thomas Müller, Syrine Krichene and Julian Eisenschlos .....	512
<i>Open-Domain Question Answering Goes Conversational via Question Rewriting</i>	
Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, Shayne Longpre, Stephen Pulman and Srinivas Chappidi .....	520
<i>QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering</i>	
Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang and Jure Leskovec .....	535
<i>XOR QA: Cross-lingual Open-Retrieval Question Answering</i>	
Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi and Hannaneh Hajishirzi ..	547
<i>SPARTA: Efficient Open-Domain Question Answering via Sparse Transformer Matching Retrieval</i>	
Tiancheng Zhao, Xiaopeng Lu and Kyusong Lee .....	565

<i>Implicitly Abusive Language – What does it actually look like and why are we not getting there?</i>	
Michael Wiegand, Josef Ruppenhofer and Elisabeth Eder .....	576
<i>The Importance of Modeling Social Factors of Language: Theory and Practice</i>	
Dirk Hovy and Diyi Yang .....	588
<i>On learning and representing social meaning in NLP: a sociolinguistic perspective</i>	
Dong Nguyen, Laura Rosseel and Jack Grieve .....	603
<i>Preregistering NLP research</i>	
Emiel van Miltenburg, Chris van der Lee and Emiel Kraemer .....	613
<i>Get Your Vitamin C! Robust Fact Verification with Contrastive Evidence</i>	
Tal Schuster, Adam Fisch and Regina Barzilay .....	624
<i>Representing Numbers in NLP: a Survey and a Vision</i>	
Avijit Thawani, Jay Pujara, Filip Ilievski and Pedro Szekely .....	644
<i>Extending Multi-Document Summarization Evaluation to the Interactive Setting</i>	
Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer and Ido Dagan	
657	
<i>Identifying Helpful Sentences in Product Reviews</i>	
Iftah Gamzu, Hila Gonen, Gilad Kutiel, Ran Levy and Eugene Agichtein .....	678
<i>Noisy Self-Knowledge Distillation for Text Summarization</i>	
Yang Liu, Sheng Shen and Mirella Lapata .....	692
<i>Improving Zero and Few-Shot Abstractive Summarization with Intermediate Fine-tuning and Data Augmentation</i>	
Alexander Fabbri, Simeng Han, Haoyuan Li, Haoran Li, Marjan Ghazvininejad, Shafiq Joty, Dragomir Radev and Yashar Mehdad .....	704
<i>Enhancing Factual Consistency of Abstractive Summarization</i>	
Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang and Meng Jiang .....	718
<i>Few-shot Intent Classification and Slot Filling with Retrieved Examples</i>	
Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat and Qi Li .....	734
<i>"Nice Try, Kiddo": Investigating Ad Hominems in Dialogue Responses</i>	
Emily Sheng, Kai-Wei Chang, Prem Natarajan and Nanyun Peng .....	750
<i>Human-like informative conversations: Better acknowledgements using conditional mutual information</i>	
Ashwin Paranjape and Christopher Manning .....	768
<i>A Comparative Study on Schema-Guided Dialogue State Tracking</i>	
Jie Cao and Yi Zhang .....	782
<i>Spoken Language Understanding for Task-oriented Dialogue Systems with Augmented Memory Networks</i>	
Jie Wu, Ian Harris and Hongzhi Zhao .....	797
<i>How to Motivate Your Dragon: Teaching Goal-Driven Agents to Speak and Act in Fantasy Worlds</i>	
Prithviraj Ammanabrolu, Jack Urbanek, Margaret Li, Arthur Szlam, Tim Rocktäschel and Jason Weston .....	807

<i>Linking Entities to Unseen Knowledge Bases with Arbitrary Schemas</i> Yogarshi Vyas and Miguel Ballesteros .....	834
<i>Self-Training with Weak Supervision</i> Giannis Karamanolakis, Subhabrata Mukherjee, Guoqing Zheng and Ahmed Hassan Awadallah	845
<i>Neural Language Modeling for Contextualized Temporal Graph Generation</i> Aman Madaan and Yiming Yang .....	864
<i>Probabilistic Box Embeddings for Uncertain Knowledge Graph Reasoning</i> Xuelu Chen, Michael Boratko, Muhao Chen, Shib Sankar Dasgupta, Xiang Lorraine Li and Andrew McCallum .....	882
<i>Document-Level Event Argument Extraction by Conditional Generation</i> Sha Li, Heng Ji and Jiawei Han .....	894
<i>Template Filling with Generative Transformers</i> Xinya Du, Alexander Rush and Claire Cardie .....	909
<i>Towards Interpreting and Mitigating Shortcut Learning Behavior of NLU models</i> Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun and Xia Hu .....	915
<i>On Attention Redundancy: A Comprehensive Study</i> Yuchen Bian, Jiaji Huang, Xingyu Cai, Jiahong Yuan and Kenneth Church .....	930
<i>Does BERT Pretrained on Clinical Notes Reveal Sensitive Data?</i> Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg and Byron Wallace .....	946
<i>Low-Complexity Probing via Finding Subnetworks</i> Victor Sanh and Alexander Rush .....	960
<i>An Empirical Comparison of Instance Attribution Methods for NLP</i> Pouya Pezeshkpour, Sarthak Jain, Byron Wallace and Sameer Singh .....	967
<i>Generalization in Instruction Following Systems</i> Soham Dan, Michael Zhou and Dan Roth .....	976
<i>LightningDOT: Pre-training Visual-Semantic Embeddings for Real-Time Image-Text Retrieval</i> Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang and Jingjing Liu .....	982
<i>Measuring Social Biases in Grounded Vision and Language Embeddings</i> Candace Ross, Boris Katz and Andrei Barbu .....	998
<i>MTAG: Modal-Temporal Attention Graph for Unaligned Human Multimodal Language Sequences</i> Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria and Louis-Philippe Morency .....	1009
<i>Grounding Open-Domain Instructions to Automate Web Support Tasks</i> Nancy Xu, Sam Masling, Michael Du, Giovanni Campagna, Larry Heck, James Landay and Monica Lam .....	1022
<i>Modular Networks for Compositional Instruction Following</i> Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein and trevor darrell .....	1033

<i>Improving Cross-Modal Alignment in Vision Language Navigation via Syntactic Information</i>	
Jialu Li, Hao Tan and Mohit Bansal .....	1041
<i>Improving Pretrained Models for Zero-shot Multi-label Text Classification through Reinforced Label Hierarchy Reasoning</i>	
Hui Liu, Danqing Zhang, Bing Yin and Xiaodan Zhu .....	1051
<i>Fine-Tuning Pre-trained Language Model with Weak Supervision: A Contrastive-Regularized Self-Training Approach</i>	
Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao and Chao Zhang .....	1063
<i>Posterior Differential Regularization with <math>f</math>-divergence for Improving Model Robustness</i>	
Hao Cheng, Xiaodong Liu, Lis Pereira, Yaoliang Yu and Jianfeng Gao .....	1078
<i>Understanding Hard Negatives in Noise Contrastive Estimation</i>	
Wenzheng Zhang and Karl Stratos .....	1090
<i>Certified Robustness to Word Substitution Attack with Differential Privacy</i>	
Wenjie Wang, Pengfei Tang, Jian Lou and Li Xiong .....	1102
<i>DReCa: A General Task Augmentation Strategy for Few-Shot Natural Language Inference</i>	
Shikhar Murty, Tatsunori Hashimoto and Christopher Manning .....	1113
<i>Harnessing Multilinguality in Unsupervised Machine Translation for Rare Languages</i>	
Xavier Garcia, Aditya Siddhant, Orhan Firat and Ankur Parikh .....	1126
<i>Macro-Average: Rare Types Are Important Too</i>	
Thamme Gowda, Weiqiu You, Constantine Lignos and Jonathan May .....	1138
<i>Assessing Reference-Free Peer Evaluation for Machine Translation</i>	
Sweta Agrawal, George Foster, Markus Freitag and Colin Cherry .....	1158
<i>The Curious Case of Hallucinations in Neural Machine Translation</i>	
Vikas Raunak, Arul Menezes and Marcin Junczys-Dowmunt .....	1172
<i>Towards Continual Learning for Multilingual Machine Translation via Vocabulary Substitution</i>	
Xavier Garcia, Noah Constant, Ankur Parikh and Orhan Firat .....	1184
<i>Towards Modeling the Style of Translators in Neural Machine Translation</i>	
Yue Wang, Cuong Hoang and Marcello Federico .....	1193
<i>Self-Supervised Test-Time Learning for Reading Comprehension</i>	
Pratyay Banerjee, Tejas Gokhale and Chitta Baral .....	1200
<i>Capturing Row and Column Semantics in Transformer Based Question Answering over Tables</i>	
Michael Glass, Mustafa Canim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avi Sil, Feifei Pan, Samarth Bharadwaj and Nicolas Rodolfo Fauceglier .....	1212
<i>Explainable Multi-hop Verbal Reasoning Through Internal Monologue</i>	
Zhengzhong Liang, Steven Bethard and Mihai Surdeanu .....	1225
<i>Robust Question Answering Through Sub-part Alignment</i>	
Jifan Chen and Greg Durrett .....	1251
<i>Text Modular Networks: Learning to Decompose Tasks in the Language of Existing Models</i>	
Tushar Khot, Daniel Khoshdel, Kyle Richardson, Peter Clark and Ashish Sabharwal .....	1264

<i>RECONSIDER: Improved Re-Ranking using Span-Focused Cross-Attention for Open Domain Question Answering</i>	
Srinivasan Iyer, Sewon Min, Yashar Mehdad and Wen-tau Yih .....	1280
<i>On the Transferability of Minimal Prediction Preserving Inputs in Question Answering</i>	
Shayne Longpre, Yi Lu and Chris DuBois .....	1288
<i>Understanding by Understanding Not: Modeling Negation in Language Models</i>	
Arian Hosseini, Siva Reddy, Dzmitry Bahdanau, R Devon Hjelm, Alessandro Sordani and Aaron Courville .....	1301
<i>DuoRAT: Towards Simpler Text-to-SQL Models</i>	
Torsten Scholak, Raymond Li, Dzmitry Bahdanau, Harm de Vries and Chris Pal .....	1313
<i>Looking Beyond Sentence-Level Natural Language Inference for Question Answering and Text Summarization</i>	
Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Lorraine Li, Pavan Kapanipathi and Kartik Talamadupula .....	1322
<i>Structure-Grounded Pretraining for Text-to-SQL</i>	
Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun and Matthew Richardson .....	1337
<i>Incremental Few-shot Text Classification with Multi-round New Classes: Formulation, Dataset and System</i>	
Congying Xia, Wenpeng Yin, Yihao Feng and Philip Yu .....	1351
<i>Temporal Reasoning on Implicit Events from Distant Supervision</i>	
Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal and Dan Roth ....	1361
<i>Disentangling Semantics and Syntax in Sentence Embeddings with Pre-trained Language Models</i>	
James Y. Huang, Kuan-Hao Huang and Kai-Wei Chang .....	1372
<i>Structure-Aware Abstractive Conversation Summarization via Discourse and Action Graphs</i>	
Jiaao Chen and Diyi Yang .....	1380
<i>A New Approach to Overgenerating and Scoring Abstractive Summaries</i>	
Kaiqiang Song, Bingqing Wang, Zhe Feng and Fei Liu .....	1392
<i>D2S: Document-to-Slide Generation Via Query-Based Text Summarization</i>	
Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang and Nancy X. R. Wang .....	1405
<i>Efficient Attentions for Long Document Summarization</i>	
Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji and Lu Wang .....	1419
<i>RefSum: Refactoring Neural Summarization</i>	
Yixin Liu, Zi-Yi Dou and Pengfei Liu .....	1437
<i>Annotating and Modeling Fine-grained Factuality in Summarization</i>	
Tanya Goyal and Greg Durrett .....	1449
<i>Larger-Context Tagging: When and Why Does It Work?</i>	
Jinlan Fu, Liangjing Feng, Qi Zhang, Xuanjing Huang and Pengfei Liu .....	1463
<i>Neural Sequence Segmentation as Determining the Leftmost Segments</i>	
Yangming Li, Lemao Liu and Kaisheng Yao .....	1476

<i>PCFGs Can Do Better: Inducing Probabilistic Context-Free Grammars with Many Symbols</i> Songlin Yang, Yanpeng Zhao and Kewei Tu .....	1487
<i>GEMNET: Effective Gated Gazetteer Representations for Recognizing Complex Entities in Low-context Input</i> Tao Meng, Anjie Fang, Oleg Rokhlenko and Shervin Malmasi .....	1499
<i>Video-aided Unsupervised Grammar Induction</i> Songyang Zhang, Linfeng Song, Lifeng Jin, Kun Xu, Dong Yu and Jiebo Luo .....	1513
<i>Generating Negative Samples by Manipulating Golden Responses for Unsupervised Learning of a Response Evaluation Model</i> ChaeHun Park, Eugene Jang, Wonsuk Yang and Jong Park .....	1525
<i>How Robust are Fact Checking Systems on Colloquial Claims?</i> Byeongchang Kim, Hyunwoo Kim, Seokhee Hong and Gunhee Kim .....	1535
<i>Fine-grained Post-training for Improving Retrieval-based Dialogue Systems</i> Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko and Jungyun Seo .....	1549
<i>Put Chatbot into Its Interlocutor's Shoes: New Framework to Learn Chatbot Responding with Intention</i> Hsuan Su, Jiun-Hao Jhan, Fan-yun Sun, Saurav Sahay and Hung-yi Lee .....	1559
<i>Adding Chit-Chat to Enhance Task-Oriented Dialogues</i> Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho and Claire Cardie .....	1570
<i>Incorporating Syntax and Semantics in Coreference Resolution with Heterogeneous Graph Attention Network</i> Fan Jiang and Trevor Cohn .....	1584
<i>Context Tracking Network: Graph-based Context Modeling for Implicit Discourse Relation Recognition</i> Yingxue Zhang, Fandong Meng, Peng Li, Ping Jian and Jie Zhou .....	1592
<i>Improving Neural RST Parsing Model with Silver Agreement Subtrees</i> Naoki Kobayashi, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura and Masaaki Nagata	1600
<i>RST Parsing from Scratch</i> Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty and Xiaoli Li .....	1613
<i>Did they answer? Subjective acts and intents in conversational discourse</i> Elisa Ferracane, Greg Durrett, Junyi Jessy Li and Katrin Erk .....	1626
<i>Evaluating the Impact of a Hierarchical Discourse Representation on Entity Coreference Resolution Performance</i> Sopan Khosla, James Fiacco and Carolyn Rosé .....	1645
<i>Bridging Resolution: Making Sense of the State of the Art</i> Hideo Kobayashi and Vincent Ng .....	1652
<i>Explicitly Modeling Syntax in Language Models with Incremental Parsing and a Dynamic Oracle</i> Yikang Shen, Shawn Tan, Alessandro Sordani, Siva Reddy and Aaron Courville .....	1660
<i>Revisiting the Weaknesses of Reinforcement Learning for Neural Machine Translation</i> Samuel Kiegeand and Julia Kreutzer .....	1673

<i>Learning to Organize a Bag of Words into Sentences with Neural Networks: An Empirical Study</i> Chongyang Tao, Shen Gao, Juntao Li, Yansong Feng, Dongyan Zhao and Rui Yan .....	1682
<i>Mask Attention Networks: Rethinking and Strengthen Transformer</i> Zhihao Fan, Yeyun Gong, Dayiheng Liu, Zhongyu Wei, Siyuan Wang, Jian Jiao, Nan Duan, Ruofei Zhang and Xuanjing Huang .....	1692
<i>ERNIE-Gram: Pre-Training with Explicitly N-Gram Masked Language Modeling for Natural Language Understanding</i> Dongling Xiao, Yu-Kun Li, Han Zhang, Yu Sun, Hao Tian, Hua Wu and Haifeng Wang .....	1702
<i>Lattice-BERT: Leveraging Multi-Granularity Representations in Chinese Pre-trained Language Models</i> Yuxuan Lai, Yijia Liu, Yansong Feng, Songfang Huang and Dongyan Zhao .....	1716
<i>Modeling Event Plausibility with Consistent Conceptual Abstraction</i> Ian Porada, Kaheer Suleman, Adam Trischler and Jackie Chi Kit Cheung .....	1732
<i>UmlsBERT: Clinical Domain Knowledge Augmentation of Contextual Embeddings Using the Unified Medical Language System Metathesaurus</i> George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen and Alexander Wong ....	1744
<i>Field Embedding: A Unified Grain-Based Framework for Word Representation</i> Junjie Luo, Xi Chen, Jichao Sun, Yuejia Xiang, Ningyu Zhang and Xiang Wan .....	1754
<i>MelBERT: Metaphor Detection via Contextualized Late Interaction using Metaphorical Identification Theories</i> Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee and Jongwuk Lee .....	1763
<i>Non-Parametric Few-Shot Learning for Word Sense Disambiguation</i> Howard Chen, Mengzhou Xia and Danqi Chen .....	1774
<i>Why Do Document-Level Polarity Classifiers Fail?</i> Karen Martins, Pedro O.S Vaz-de-Melo and Rodrygo Santos .....	1782
<i>A Unified Span-Based Approach for Opinion Mining with Syntactic Constituents</i> Qingrong Xia, Bo Zhang, Rui Wang, Zhenghua Li, Yue Zhang, Fei Huang, Luo Si and Min Zhang	1795
<i>Target-specified Sequence Labeling with Multi-head Self-attention for Target-oriented Opinion Words Extraction</i> Yuhao Feng, Yanghui Rao, Yuyao Tang, Ninghua Wang and He Liu .....	1805
<i>Does syntax matter? A strong baseline for Aspect-based Sentiment Analysis with RoBERTa</i> Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu and Xipeng Qiu .....	1816
<i>Domain Divergences: A Survey and Empirical Analysis</i> Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan and Roger Zimmermann ...	1830
<i>Target-Aware Data Augmentation for Stance Detection</i> Yingjie Li and Cornelia Caragea .....	1850
<i>End-to-end ASR to jointly predict transcriptions and linguistic annotations</i> Motoi Omachi, Yuya Fujita, Shinji Watanabe and Matthew Wiesner .....	1861

<i>Source and Target Bidirectional Knowledge Distillation for End-to-end Speech Translation</i>	
Hirofumi Inaguma, Tatsuya Kawahara and Shinji Watanabe .....	1872
<i>Searchable Hidden Intermediates for End-to-End Models of Decomposable Sequence Tasks</i>	
Siddharth Dalmia, Brian Yan, Vikas Raunak, Florian Metze and Shinji Watanabe .....	1882
<i>SPLAT: Speech-Language Joint Pre-Training for Spoken Language Understanding</i>	
Yu-An Chung, Chenguang Zhu and Michael Zeng .....	1897
<i>Worldly Wise (WoW) - Cross-Lingual Knowledge Fusion for Fact-based Visual Spoken-Question Answering</i>	
Kiran Ramnath, Leda Sari, Mark Hasegawa-Johnson and Chang Yoo .....	1908
<i>Align-Refine: Non-Autoregressive Speech Recognition via Iterative Realignment</i>	
Ethan A. Chi, Julian Salazar and Katrin Kirchhoff .....	1920
<i>Everything Has a Cause: Leveraging Causal Inference in Legal Text Analysis</i>	
Xiao Liu, Da Yin, Yansong Feng, Yuting Wu and Dongyan Zhao .....	1928
<i>Counterfactual Supporting Facts Extraction for Explainable Medical Record Based Diagnosis with Graph Network</i>	
Haoran Wu, Wei Chen, Shuang Xu and Bo Xu .....	1942
<i>Personalized Response Generation via Generative Split Memory Network</i>	
Yuwei Wu, Xuezhe Ma and Diyi Yang .....	1956
<i>Towards Few-shot Fact-Checking via Perplexity</i>	
Nayeon Lee, Yejin Bang, Andrea Madotto and Pascale Fung .....	1971
<i>Active<sup>2</sup> Learning: Actively reducing redundancies in Active Learning methods for Sequence Tagging and Machine Translation</i>	
Rishi Hazra, Parag Dutta, Shubham Gupta, Mohammed Abdul Qaathir and Ambedkar Dukkipati	1982
<i>Generating An Optimal Interview Question Plan Using A Knowledge Graph And Integer Linear Programming</i>	
Soham Datta, Prabir Mallick, Sangameshwar Patil, Indrajit Bhattacharya and Girish Palshikar	1996
<i>Model Extraction and Adversarial Transferability, Your BERT is Vulnerable!</i>	
Xuanli He, Lingjuan Lyu, Lichao Sun and Qionгкаi Xu .....	2006
<i>A Global Past-Future Early Exit Method for Accelerating Inference of Pre-trained Language Models</i>	
Kaiyuan Liao, Yi Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He .....	2013
<i>Masked Conditional Random Fields for Sequence Labeling</i>	
Tianwen Wei, Jianwei Qi, Shenghuan He and Songtao Sun .....	2024
<i>Heterogeneous Graph Neural Networks for Concept Prerequisite Relation Learning in Educational Data</i>	
Chenghao Jia, Yongliang Shen, Yechun Tang, Lu Sun and Weiming Lu .....	2036
<i>Be Careful about Poisoned Word Embeddings: Exploring the Vulnerability of the Embedding Layers in NLP Models</i>	
Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun and Bin He .....	2048
<i>DA-Transformer: Distance-aware Transformer</i>	
Chuhan Wu, Fangzhao Wu and Yongfeng Huang .....	2059



<i>ASAP: A Chinese Review Dataset Towards Aspect Category Sentiment Analysis and Rating Prediction</i> Jiahao Bu, Lei Ren, Shuang Zheng, Yang Yang, Jingang Wang, Fuzheng Zhang and Wei Wu .	2069
<i>Are NLP Models really able to Solve Simple Math Word Problems?</i> Arkil Patel, Satwik Bhattamishra and Navin Goyal . . . . .	2080
<i>WRIME: A New Dataset for Emotional Intensity Estimation with Subjective and Objective Annotations</i> Tomoyuki Kajiwara, Chenhui Chu, Noriko Takemura, Yuta Nakashima and Hajime Nagahara	2095
<i>KPQA: A Metric for Generative Question Answering Using Keyphrase Weights</i> Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Joongbo Shin and Kyomin Jung . . . . .	2105
<i>StylePTB: A Compositional Benchmark for Fine-grained Controllable Text Style Transfer</i> Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov and Louis-Philippe Morency . . . . .	2116
<i>Blow the Dog Whistle: A Chinese Dataset for Cant Understanding with Common Sense and World Knowledge</i> Canwen Xu, Wangchunshu Zhou, Tao Ge, Ke Xu, Julian McAuley and Furu Wei . . . . .	2139
<i>COVID-19 Named Entity Recognition for Vietnamese</i> Thinh Hung Truong, Mai Hoang Dao and Dat Quoc Nguyen . . . . .	2146
<i>Framing Unpacked: A Semi-Supervised Interpretable Multi-View Model of Media Frames</i> Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak, Andrew Turpin and Lea Frermann . . . . .	2154
<i>Automatic Classification of Neutralization Techniques in the Narrative of Climate Change Scepticism</i> Shraey Bhatia, Jey Han Lau and Timothy Baldwin . . . . .	2167
<i>Suicide Ideation Detection via Social and Temporal User Representations using Hyperbolic Learning</i> Ramit Sawhney, Harshit Joshi, Rajiv Ratn Shah and Lucie Flek . . . . .	2176
<i>WikiTalkEdit: A Dataset for modeling Editors' behaviors on Wikipedia</i> Kokil Jaidka, Andrea Ceolin, Iknor Singh, Niyati Chhaya and Lyle Ungar . . . . .	2191
<i>The structure of online social networks modulates the rate of lexical change</i> Jian Zhu and David Jurgens . . . . .	2201
<i>Modeling Framing in Immigration Discourse on Social Media</i> Julia Mendelsohn, Ceren Budak and David Jurgens . . . . .	2219
<i>Modeling the Severity of Complaints in Social Media</i> Mali Jin and Nikolaos Aletras . . . . .	2264
<i>What About the Precedent: An Information-Theoretic Analysis of Common Law</i> Josef Valvoda, Tiago Pimentel, Niklas Stoehr, Ryan Cotterell and Simone Teufel . . . . .	2275
<i>Introducing CAD: the Contextual Abuse Dataset</i> Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini and Rebekah Tromble . . . . .	2289
<i>Lifelong Learning of Hate Speech Classification on Social Media</i> Jing Qian, Hong Wang, Mai ElSherief and Xifeng Yan . . . . .	2304

<i>Learning to Recognize Dialect Features</i>	
Dorottya Demszky, Devyani Sharma, Jonathan Clark, Vinodkumar Prabhakaran and Jacob Eisenstein .....	2315
<i>It's Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners</i>	
Timo Schick and Hinrich Schütze .....	2339
<i>Static Embeddings as Efficient Knowledge Bases?</i>	
Philipp Dufter, Nora Kassner and Hinrich Schütze .....	2353
<i>Highly Efficient Knowledge Graph Embedding Learning with Orthogonal Procrustes Analysis</i>	
Xutan Peng, Guanyi Chen, Chenghua Lin and Mark Stevenson .....	2364
<i>Rethinking Network Pruning – under the Pre-train and Fine-tune Paradigm</i>	
Dongkuan Xu, Ian En-Hsu Yen, Jinxi Zhao and Zhibin Xiao .....	2376
<i>Towards a Comprehensive Understanding and Accurate Evaluation of Societal Biases in Pre-Trained Transformers</i>	
Andrew Silva, Pradyumna Tambwekar and Matthew Gombolay .....	2383
<i>Detoxifying Language Models Risks Marginalizing Minority Voices</i>	
Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap and Dan Klein ...	2390
<i>HONEST: Measuring Hurtful Sentence Completion in Language Models</i>	
Debora Nozza, Federico Bianchi and Dirk Hovy .....	2398
<i>EaSe: A Diagnostic Tool for VQA based on Answer Diversity</i>	
Shailza Jolly, Sandro Pezzelle and Moin Nabi .....	2407
<i>DeCEMBERT: Learning from Noisy Instructional Videos via Dense Captions and Entropy Minimization</i>	
Zineng Tang, Jie Lei and Mohit Bansal .....	2415
<i>Improving Generation and Evaluation of Visual Stories via Semantic Consistency</i>	
Adyasha Maharana, Darryl Hannan and Mohit Bansal .....	2427
<i>Multilingual Multimodal Pre-training for Zero-Shot Cross-Lingual Transfer of Vision-Language Models</i>	
Po-Yao Huang, Mandela Patrick, Junjie Hu, Graham Neubig, Florian Metze and Alexander Hauptmann .....	2443
<i>Video Question Answering with Phrases via Semantic Roles</i>	
Arka Sadhu, Kan Chen and Ram Nevatia .....	2460
<i>From Masked Language Modeling to Translation: Non-English Auxiliary Tasks Improve Zero-shot Spoken Language Understanding</i>	
Rob van der Goot, Ibrahim Sharaf, Aizhan Imankulova, Ahmet Üstün, Marija Stepanović, Alan Ramponi, Siti Oryza Khairunnisa, Mamoru Komachi and Barbara Plank .....	2479
<i>WEC: Deriving a Large-scale Cross-document Event Coreference dataset from Wikipedia</i>	
Alon Eirew, Arie Cattan and Ido Dagan .....	2498
<i>Challenging distributional models with a conceptual network of philosophical terms</i>	
Yvette Oortwijn, Jelke Bloem, Pia Sommerauer, Francois Meyer, Wei Zhou and Antske Fokkens	2511

<i>KILT: a Benchmark for Knowledge Intensive Language Tasks</i>	
Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel and Sebastian Riedel . . . . .	2523
<i>A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios</i>	
Michael A. Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen and Dietrich Klakow . . . . .	2545
<i>Temporal Knowledge Graph Completion using a Linear Temporal Regularizer and Multivector Embeddings</i>	
Chengjin Xu, Yung-Yu Chen, Mojtaba Nayeri and Jens Lehmann . . . . .	2569
<i>UDALM: Unsupervised Domain Adaptation through Language Modeling</i>	
Constantinos Karouzou, Georgios Paraskevopoulos and Alexandros Potamianos . . . . .	2579
<i>Beyond Black &amp; White: Leveraging Annotator Disagreement via Soft-Label Multi-Task Learning</i>	
Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy and Massimo Poesio	2591
<i>Clustering-based Inference for Biomedical Entity Linking</i>	
Rico Angell, Nicholas Monath, Sunil Mohan, Nishant Yadav and Andrew McCallum . . . . .	2598
<i>Variance-reduced First-order Meta-learning for Natural Language Processing Tasks</i>	
Lingxiao Wang, Kevin Huang, Tengyu Ma, Quanquan Gu and Jing Huang . . . . .	2609
<i>Diversity-Aware Batch Active Learning for Dependency Parsing</i>	
Tianze Shi, Adrian Benton, Igor Malioutov and Ozan İrsoy . . . . .	2616
<i>How many data points is a prompt worth?</i>	
Teven Le Scao and Alexander Rush . . . . .	2627
<i>Can Latent Alignments Improve Autoregressive Machine Translation?</i>	
Adi Haviv, Lior Vassertail and Omer Levy . . . . .	2637
<i>Smoothing and Shrinking the Sparse Seq2Seq Search Space</i>	
Ben Peters and André F. T. Martins . . . . .	2642
<i>Unified Pre-training for Program Understanding and Generation</i>	
Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray and Kai-Wei Chang . . . . .	2655
<i>Hyperparameter-free Continuous Learning for Domain Classification in Natural Language Understanding</i>	
Ting Hua, Yilin Shen, Changsheng Zhao, Yen-Chang Hsu and Hongxia Jin . . . . .	2669
<i>On the Embeddings of Variables in Recurrent Neural Networks for Source Code</i>	
Nadezhda Chirkova . . . . .	2679
<i>Cross-Lingual Word Embedding Refinement by <math>\ell_1</math> Norm Optimisation</i>	
Xutan Peng, Chenghua Lin and Mark Stevenson . . . . .	2690
<i>Semantic Frame Forecast</i>	
Chieh-Yang Huang and Ting-Hao Huang . . . . .	2702
<i>MUSER: MULTimodal Stress detection using Emotion Recognition as an Auxiliary Task</i>	
Yiqun Yao, Michalis Papakostas, Mihai Burzo, Mohamed Abouelenien and Rada Mihalcea . .	2714

<i>Learning to Decompose and Organize Complex Tasks</i>	
Yi Zhang, Sujay Kumar Jauhar, Julia Kiseleva, Ryan White and Dan Roth . . . . .	2726
<i>Continual Learning for Text Classification with Information Disentanglement Based Regularization</i>	
Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang and Diyi Yang . . . . .	2736
<i>Learning from Executions for Semantic Parsing</i>	
Bailin Wang, Mirella Lapata and Ivan Titov . . . . .	2747
<i>Learning to Synthesize Data for Semantic Parsing</i>	
Bailin Wang, Wenpeng Yin, Xi Victoria Lin and Caiming Xiong . . . . .	2760
<i>Edge: Enriching Knowledge Graph Embeddings with External Text</i>	
Saed Rezaei, Handong Zhao, Sungchul Kim, Ryan Rossi, Nedim Lipka and Sheng Li . . . . .	2767
<i>FLIN: A Flexible Natural Language Interface for Web Navigation</i>	
Sahisnu Mazumder and Oriana Riva . . . . .	2777
<i>Game-theoretic Vocabulary Selection via the Shapley Value and Banzhaf Index</i>	
Roma Patel, Marta Garnelo, Ian Gemp, Chris Dyer and Yoram Bachrach . . . . .	2789
<i>Incorporating External Knowledge to Enhance Tabular Reasoning</i>	
J. Neeraja, Vivek Gupta and Vivek Srikumar . . . . .	2799
<i>Compositional Generalization for Neural Semantic Parsing via Span-level Supervised Attention</i>	
Pengcheng Yin, Hao Fang, Graham Neubig, Adam Pauls, Emmanouil Antonios Platanios, Yu Su, Sam Thomson and Jacob Andreas . . . . .	2810
<i>Domain Adaptation for Arabic Cross-Domain and Cross-Dialect Sentiment Analysis from Contextualized Word Embedding</i>	
Abdellah El Mekki, Abdelkader El Mahdaouy, Ismail Berrada and Ahmed Khoumsi . . . . .	2824
<i>Multi-task Learning of Negation and Speculation for Targeted Sentiment Classification</i>	
Andrew Moore and Jeremy Barnes . . . . .	2838
<i>A Disentangled Adversarial Neural Topic Model for Separating Opinions from Plots in User Reviews</i>	
Gabriele Pergola, Lin Gui and Yulan He . . . . .	2870
<i>Graph Ensemble Learning over Multiple Dependency Trees for Aspect-level Sentiment Classification</i>	
Xiao Chen Hou, Peng Qi, Guangtao Wang, Rex Ying, Jing Huang, Xiaodong He and Bowen Zhou	2884
<i>Emotion-Infused Models for Explainable Psychological Stress Detection</i>	
Elsbeth Turcan, Smaranda Muresan and Kathleen McKeown . . . . .	2895
<i>Aspect-based Sentiment Analysis with Type-aware Graph Convolutional Networks and Layer Ensemble</i>	
Yuanhe Tian, Guimin Chen and Yan Song . . . . .	2910
<i>Supertagging-based Parsing with Linear Context-free Rewriting Systems</i>	
Thomas Ruprecht and Richard Mörbitz . . . . .	2923
<i>Outside Computation with Superior Functions</i>	
Parker Riley and Daniel Gildea . . . . .	2936
<i>Learning Syntax from Naturally-Occurring Bracketings</i>	
Tianze Shi, Ozan Irsoy, Igor Malioutov and Lillian Lee . . . . .	2941

<i>Bot-Adversarial Dialogue for Safe Conversational Agents</i>	
Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston and Emily Dinan . . . . .	2950
<i>Non-Autoregressive Semantic Parsing for Compositional Task-Oriented Dialog</i>	
Arun Babu, Akshat Shrivastava, Armen Aghajanyan, Ahmed Aly, Angela Fan and Marjan Ghazvininejad . . . . .	2969
<i>Example-Driven Intent Prediction with Observers</i>	
Shikib Mehri and Mihail Eric . . . . .	2979
<i>Imperfect also Deserves Reward: Multi-Level and Sequential Reward Modeling for Better Dialog Management</i>	
Zhengxu Hou, Bang Liu, Ruihui Zhao, Zijing Ou, Yafei Liu, Xi Chen and Yefeng Zheng . . . . .	2993
<i>Action-Based Conversations Dataset: A Corpus for Building More In-Depth Task-Oriented Dialogue Systems</i>	
Derek Chen, Howard Chen, Yi Yang, Alexander Lin and Zhou Yu . . . . .	3002
<i>Controlling Dialogue Generation with Semantic Exemplars</i>	
Prakhar Gupta, Jeffrey Bigham, Yulia Tsvetkov and Amy Pavel . . . . .	3018
<i>COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List</i>	
Luyu Gao, Zhuyun Dai and Jamie Callan . . . . .	3030
<i>X-Class: Text Classification with Extremely Weak Supervision</i>	
Zihan Wang, Dheeraj Mekala and Jingbo Shang . . . . .	3043
<i>Fine-tuning Encoders for Improved Monolingual and Zero-shot Polylingual Neural Topic Modeling</i>	
Aaron Mueller and Mark Dredze . . . . .	3054
<i>Exploring the Relationship Between Algorithm Performance, Vocabulary, and Run-Time in Text Classification</i>	
Wilson Fearn, Orion Weller and Kevin Seppi . . . . .	3069
<i>Faithfully Explainable Recommendation via Neural Logic Reasoning</i>	
Yaxin Zhu, Yikun Xian, Zuohui Fu, Gerard de Melo and Yongfeng Zhang . . . . .	3083
<i>You Sound Like Someone Who Watches Drama Movies: Towards Predicting Movie Preferences from Conversational Interactions</i>	
Sergey Volokhin, Joyce Ho, Oleg Rokhlenko and Eugene Agichtein . . . . .	3091
<i>Reading and Acting while Blindfolded: The Need for Semantics in Text Game Agents</i>	
Shunyu Yao, Karthik Narasimhan and Matthew Hausknecht . . . . .	3097
<i>SOrT-ing VQA Models : Contrastive Gradient Learning for Improved Consistency</i>	
Sameer Dharur, Purva Tendulkar, Dhruv Batra, Devi Parikh and Ramprasaath R. Selvaraju . . . . .	3103
<i>Semi-Supervised Policy Initialization for Playing Games with Language Hints</i>	
Tsu-Jui Fu and William Yang Wang . . . . .	3112
<i>Revisiting Document Representations for Large-Scale Zero-Shot Learning</i>	
Jihyung Kil and Wei-Lun Chao . . . . .	3117
<i>Negative language transfer in learner English: A new dataset</i>	
Leticia Farias Wanderley, Nicole Zhao and Carrie Demmans Epp . . . . .	3129

<i>SentSim: Crosslingual Semantic Evaluation of Machine Translation</i> Yurun Song, Junchen Zhao and Lucia Specia .....	3143
<i>Quality Estimation for Image Captions Based on Large-scale Human Evaluations</i> Tomer Levinboim, Ashish V. Thapliyal, Piyush Sharma and Radu Soricut .....	3157
<i>CaSiNo: A Corpus of Campsite Negotiation Dialogues for Automatic Negotiation Systems</i> Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May and Jonathan Gratch	3167
<i>News Headline Grouping as a Challenging NLU Task</i> Philippe Laban, Lucas Bandarkar and Marti A. Hearst .....	3186
<i>Olá, Bonjour, Salve! XFORMAL: A Benchmark for Multilingual Formality Style Transfer</i> Eleftheria Briakou, Di Lu, Ke Zhang and Joel Tetreault .....	3199
<i>Grouping Words with Semantic Diversity</i> Karine Chubarian, Abdul Rafae Khan, Anastasios Sidiropoulos and Jia Xu .....	3217
<i>Noise Stability Regularization for Improving BERT Fine-tuning</i> Hang Hua, Xingjian Li, Dejing Dou, Chengzhong Xu and Jiebo Luo .....	3229
<i>FlowPrior: Learning Expressive Priors for Latent Variable Sentence Models</i> Xiaoan Ding and Kevin Gimpel .....	3242
<i>HTCInfoMax: A Global Model for Hierarchical Text Classification via Information Maximization</i> Zhongfen Deng, Hao Peng, Dongxiao He, Jianxin Li and Philip Yu .....	3259
<i>Knowledge Guided Metric Learning for Few-Shot Text Classification</i> Dianbo Sui, Yubo Chen, Binjie Mao, Delai Qiu, Kang Liu and Jun Zhao .....	3266
<i>Ensemble of MRR and NDCG models for Visual Dialog</i> Idan Schwartz .....	3272
<i>Supervised Neural Clustering via Latent Structured Output Learning: Application to Question Intents</i> Iryna Haponchyk and Alessandro Moschitti .....	3364
<i>ConVEx: Data-Efficient and Few-Shot Slot Labeling</i> Matthew Henderson and Ivan Vulić .....	3375
<i>CREAD: Combined Resolution of Ellipses and Anaphora in Dialogues</i> Bo-Hsiang Tseng, Shruti Bhargava, Jiarui Lu, Joel Ruben Antony Moniz, Dhivya Piraviperumal, Lin Li and Hong Yu .....	3390
<i>Knowledge-Driven Slot Constraints for Goal-Oriented Dialogue Systems</i> Piyawat Lertvittayakumjorn, Daniele Bonadiman and Saab Mansour .....	3407
<i>Clipping Loops for Sample-Efficient Dialogue Policy Optimisation</i> Yen-Chen Wu and Carl Edward Rasmussen .....	3420
<i>Integrating Lexical Information into Entity Neighbourhood Representations for Relation Prediction</i> Ian Wood, Mark Johnson and Stephen Wan .....	3429
<i>Noisy-Labeled NER with Confidence Estimation</i> Kun Liu, Yao Fu, Chuanqi Tan, Moshu Chen, Ningyu Zhang, Songfang Huang and Sheng Gao	3437

<i>TABBIE: Pretrained Representations of Tabular Data</i> Hiroshi Iida, Dung Thai, Varun Manjunatha and Mohit Iyyer .....	3446
<i>Better Feature Integration for Named Entity Recognition</i> Lu Xu, Zhanming Jie, Wei Lu and Lidong Bing .....	3457
<i>ZS-BERT: Towards Zero-Shot Relation Extraction with Attribute Representation Learning</i> Chih-Yao Chen and Cheng-Te Li .....	3470
<i>Graph Convolutional Networks for Event Causality Identification with Rich Document-level Structures</i> Minh Tran Phu and Thien Huu Nguyen .....	3480
<i>A Context-Dependent Gated Module for Incorporating Symbolic Semantics into Event Coreference Resolution</i> Tuan Lai, Heng Ji, Trung Bui, Quan Hung Tran, Franck Dernoncourt and Walter Chang .....	3491
<i>Multi-Style Transfer with Discriminative Feedback on Disjoint Corpus</i> Navita Goyal, Balaji Vasan Srinivasan, Anandhavelu N and Abhilasha Sancheti .....	3500
<i>FUDGE: Controlled Text Generation With Future Discriminators</i> Kevin Yang and Dan Klein .....	3511
<i>Controllable Text Simplification with Explicit Paraphrasing</i> Mounica Maddela, Fernando Alva-Manchego and Wei Xu .....	3536
<i>Knowledge Graph Based Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-training</i> Oshin Agarwal, Heming Ge, Siamak Shakeri and Rami Al-Rfou .....	3554
<i>Choose Your Own Adventure: Paired Suggestions in Collaborative Writing for Evaluating Story Generation Models</i> Elizabeth Clark and Noah A. Smith .....	3566
<i>InfoXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training</i> Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang and Ming Zhou .....	3576
<i>Context-Interactive Pre-Training for Document Machine Translation</i> Pengcheng Yang, Pei Zhang, Boxing Chen, Jun Xie and Weihua Luo .....	3589
<i>Code-Mixing on Sesame Street: Dawn of the Adversarial Polyglots</i> Samson Tan and Shafiq Joty .....	3596
<i>X-METRA-ADA: Cross-lingual Meta-Transfer learning Adaptation to Natural Language Understanding and Question Answering</i> Meryem M’hamdi, Doo Soon Kim, Franck Dernoncourt, Trung Bui, Xiang Ren and Jonathan May	3617
<i>Explicit Alignment Objectives for Multilingual Bidirectional Encoders</i> Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant and Graham Neubig .....	3633
<i>Cross-lingual Cross-modal Pretraining for Multimodal Retrieval</i> Hongliang Fei, Tan Yu and Ping Li .....	3644

<i>Wikipedia Entities as Rendezvous across Languages: Grounding Multilingual Language Models by Predicting Wikipedia Hyperlinks</i>	
Iacer Calixto, Alessandro Raganato and Tommaso Pasini . . . . .	3651
<i>multiPProver: Generating Multiple Proofs for Improved Interpretability in Rule Reasoning</i>	
Swarnadeep Saha, Prateek Yadav and Mohit Bansal . . . . .	3662
<i>Adaptable and Interpretable Neural Memory Over Symbolic Knowledge</i>	
Pat Verga, Haitian Sun, Livio Baldini Soares and William Cohen . . . . .	3678
<i>CLEVR_HYP: A Challenge Dataset and Baselines for Visual Question Answering with Hypothetical Actions over Images</i>	
Shailaja Keyur Sampat, Akshay Kumar, Yezhou Yang and Chitta Baral . . . . .	3692
<i>Refining Targeted Syntactic Evaluation of Language Models</i>	
Benjamin Newman, Kai-Siang Ang, Julia Gong and John Hewitt . . . . .	3710
<i>Universal Adversarial Attacks with Natural Triggers for Text Classification</i>	
Liwei Song, Xinwei Yu, Hsuan-Tung Peng and Karthik Narasimhan . . . . .	3724
<i>QuadrupletBERT: An Efficient Model For Embedding-Based Large-Scale Retrieval</i>	
Peiyang Liu, Sen Wang, Xi Wang, Wei Ye and Shikun Zhang . . . . .	3734
<i>Dynamically Disentangling Social Bias from Task-Oriented Representations with Adversarial Attack</i>	
Liwen Wang, Yuanmeng Yan, Keqing He, Yanan Wu and Weiran Xu . . . . .	3740
<i>An Empirical Investigation of Bias in the Multimodal Analysis of Financial Earnings Calls</i>	
Ramit Sawhney, Arshiya Aggarwal and Rajiv Ratn Shah . . . . .	3751
<i>Beyond Fair Pay: Ethical Implications of NLP Crowdsourcing</i>	
Boaz Shmueli, Jan Fell, Soumya Ray and Lun-Wei Ku . . . . .	3758
<i>On Transferability of Bias Mitigation Effects in Language Model Fine-Tuning</i>	
Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves and Xiang Ren . . . . .	3770
<i>Case Study: Deontological Ethics in NLP</i>	
Shrimai Prabhumoye, Brendon Boldt, Ruslan Salakhutdinov and Alan W Black . . . . .	3784
<i>Privacy Regularization: Joint Privacy-Utility Optimization in Language Models</i>	
Fatemehsadat Mireshghallah, Huseyin Inan, Marcello Hasegawa, Victor Rühle, Taylor Berg-Kirkpatrick and Robert Sim . . . . .	3799
<i>On the Impact of Random Seeds on the Fairness of Clinical Classifiers</i>	
Silvio Amir, Jan-Willem van de Meent and Byron Wallace . . . . .	3808
<i>Topic Model or Topic Twaddle? Re-evaluating Semantic Interpretability Measures</i>	
Caitlin Doogan and Wray Buntine . . . . .	3824
<i>Discourse Probing of Pretrained Language Models</i>	
Fajri Koto, Jey Han Lau and Timothy Baldwin . . . . .	3849
<i>UniDrop: A Simple yet Effective Technique to Improve Transformer without Extra Cost</i>	
Zhen Wu, Lijun Wu, Qi Meng, Yingce Xia, Shufang Xie, Tao Qin, Xinyu Dai and Tie-Yan Liu	3865



<i>tWT-WT: A Dataset to Assert the Role of Target Entities for Detecting Stance of Tweets</i> Ayush Kaushal, Avirup Saha and Niloy Ganguly .....	3879
<i>Learning to Learn to be Right for the Right Reasons</i> Pride Kavumba, Benjamin Heinzerling, Ana Brassard and Kentaro Inui .....	3890
<i>Double Perturbation: On the Robustness of Robustness and Counterfactual Bias Evaluation</i> Chong Zhang, Jieyu Zhao, Huan Zhang, Kai-Wei Chang and Cho-Jui Hsieh .....	3899
<i>Explaining Neural Network Predictions on Sentence Pairs via Learning Word-Group Masks</i> Hanjie Chen, Song Feng, Jatin Ganhotra, Hui Wan, Chulaka Gunasekara, Sachindra Joshi and Yangfeng Ji .....	3917
<i>Almost Free Semantic Draft for Neural Machine Translation</i> Xi Ai and Bin Fang .....	3931
<i>Pruning-then-Expanding Model for Domain Adaptation of Neural Machine Translation</i> Shuhao Gu, Yang Feng and Wanying Xie .....	3942
<i>Multi-Hop Transformer for Document-Level Machine Translation</i> Long Zhang, Tong Zhang, Haibo Zhang, Baosong Yang, Wei Ye and Shikun Zhang .....	3953
<i>Continual Learning for Neural Machine Translation</i> Yue Cao, Hao-Ran Wei, Boxing Chen and Xiaojun Wan .....	3964
<i>Self-Training for Unsupervised Neural Machine Translation in Unbalanced Training Data Scenarios</i> Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita and Tiejun Zhao .....	3975
<i>Smart-Start Decoding for Neural Machine Translation</i> Jian Yang, Shuming Ma, Dongdong Zhang, Juncheng Wan, Zhoujun Li and Ming Zhou .....	3982
<i>Multi-Task Learning with Shared Encoder for Non-Autoregressive Machine Translation</i> Yongchang Hao, Shilin He, Wenxiang Jiao, Zhaopeng Tu, Michael Lyu and Xing Wang .....	3989
<i>ER-AE: Differentially Private Text Generation for Authorship Anonymization</i> Haohan Bo, Steven H. H. Ding, Benjamin C. M. Fung and Farkhund Iqbal .....	3997
<i>Distantly Supervised Transformers For E-Commerce Product QA</i> Happy Mittal, Aniket Chakrabarti, Belhassen Bayar, Animesh Anant Sharma and Nikhil Rasiwasia	4008
<i>Quantitative Day Trading from Natural Language using Reinforcement Learning</i> Ramit Sawhney, Arnav Wadhwa, Shivam Agarwal and Rajiv Ratn Shah .....	4018
<i>Restoring and Mining the Records of the Joseon Dynasty via Neural Language Modeling and Machine Translation</i> Kyeongpil Kang, Kyohoon Jin, Soyoung Yang, Soojin Jang, Jaegul Choo and Youngbin Kim	4031
<i>Modeling Diagnostic Label Correlation for Automatic ICD Coding</i> Shang-Chi Tsai, Chao-Wei Huang and Yun-Nung Chen .....	4043
<i>Self-Supervised Contrastive Learning for Efficient User Satisfaction Prediction in Conversational Agents</i> Mohammad Kachuee, Hao Yuan, Young-Bum Kim and Sungjin Lee .....	4053
<i>A recipe for annotating grounded clarifications</i> Luciana Benotti and Patrick Blackburn .....	4065

<i>Grey-box Adversarial Attack And Defence For Sentiment Classification</i>	
Ying Xu, Xu Zhong, Antonio Jimeno Yepes and Jey Han Lau .....	4078
<i>How low is too low? A monolingual take on lemmatisation in Indian languages</i>	
Kumar Saunack, Kumar Saurav and Pushpak Bhattacharyya .....	4088
<i>Causal Effects of Linguistic Properties</i>	
Reid Pryzant, Dallas Card, Dan Jurafsky, Victor Veitch and Dhanya Sridhar .....	4095
<i>Dynabench: Rethinking Benchmarking in NLP</i>	
Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts and Adina Williams	4110
<i>Translational NLP: A New Paradigm and General Principles for Natural Language Processing Research</i>	
Denis Newman-Griffis, Jill Fain Lehman, Carolyn Rosé and Harry Hochheiser .....	4125
<i>Predicting Discourse Trees from Transformer-based Neural Summarizers</i>	
Wen Xiao, Patrick Huber and Giuseppe Carenini .....	4139
<i>Probing for Bridging Inference in Transformer Language Models</i>	
Onkar Pandit and Yufang Hou .....	4153
<i>Is Incoherence Surprising? Targeted Evaluation of Coherence Prediction from Language Models</i>	
Anne Beyer, Sharid Loáiciga and David Schlangen .....	4164
<i>Stay Together: A System for Single and Split-antecedent Anaphora Resolution</i>	
Juntao Yu, Nafise Sadat Moosavi, Silviu Paun and Massimo Poesio .....	4174
<i>Redefining Absent Keyphrases and their Effect on Retrieval Effectiveness</i>	
Florian Boudin and Ygor Gallina .....	4185
<i>CoRT: Complementary Rankings from Transformers</i>	
Marco Wrzalik and Dirk Krechel .....	4194
<i>Multi-source Neural Topic Modeling in Multi-view Embedding Spaces</i>	
Pankaj Gupta, Yatin Chaudhary and Hinrich Schütze .....	4205
<i>Inductive Topic Variational Graph Auto-Encoder for Text Classification</i>	
Qianqian Xie, Jimin Huang, Pan Du, Min Peng and Jian-Yun Nie .....	4218
<i>Self-Alignment Pretraining for Biomedical Entity Representations</i>	
Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella and Nigel Collier .....	4228
<i>TaxoClass: Hierarchical Multi-Label Text Classification Using Only Class Names</i>	
Jiaming Shen, Wenda Qiu, Yu Meng, Jingbo Shang, Xiang Ren and Jiawei Han .....	4239
<i>MERMAID: Metaphor Generation with Symbolism and Discriminative Decoding</i>	
Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan and Nanyun Peng .....	4250
<i>On Learning Text Style Transfer with Direct Rewards</i>	
Yixin Liu, Graham Neubig and John Wieting .....	4262
<i>Focused Attention Improves Document-Grounded Generation</i>	
Shrimai Prabhunoye, Kazuma Hashimoto, Yingbo Zhou, Alan W Black and Ruslan Salakhutdinov	4274

<i>NeuroLogic Decoding: (Un)supervised Neural Text Generation with Predicate Logic Constraints</i> Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula and Yejin Choi	4288
<i>Ask what’s missing and what’s useful: Improving Clarification Question Generation using Global Knowledge</i> Bodhisattwa Prasad Majumder, Sudha Rao, Michel Galley and Julian McAuley	4300
<i>Progressive Generation of Long Text with Pretrained Language Models</i> Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric Xing and Zhiting Hu	4313
<i>SOC CER: An Information-Sparse Discourse State Tracking Collection in the Sports Commentary Domain</i> Ruochen Zhang and Carsten Eickhoff	4325
<i>Plot-guided Adversarial Example Construction for Evaluating Open-domain Story Generation</i> Sarik Ghazarian, Zixi Liu, Akash S M, Ralph Weischedel, Aram Galstyan and Nanyun Peng	4334
<i>MultiOpEd: A Corpus of Multi-Perspective News Editorials</i> Siyi Liu, Sihao Chen, Xander Uyttendaele and Dan Roth	4345
<i>Swords: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality</i> Mina Lee, Chris Donahue, Robin Jia, Alexander Iyabor and Percy Liang	4362
<i>"I'm Not Mad": Commonsense Implications of Negation and Contradiction</i> Liwei Jiang, Antoine Bosselut, Chandra Bhagavatula and Yejin Choi	4380
<i>Identifying Medical Self-Disclosure in Online Communities</i> Mina Valizadeh, Pardis Ranjbar-Noiey, Cornelia Caragea and Natalie Parde	4398
<i>Language in a (Search) Box: Grounding Language Learning in Real-World Human-Machine Interaction</i> Federico Bianchi, Ciro Greco and Jacopo Tagliabue	4409
<i>Finding Concept-specific Biases in Form–Meaning Associations</i> Tiago Pimentel, Brian Roark, Søren Wichmann, Ryan Cotterell and Damián Blasi	4416
<i>How (Non-)Optimal is the Lexicon?</i> Tiago Pimentel, Irene Nikkarinen, Kyle Mahowald, Ryan Cotterell and Damián Blasi	4426
<i>Word Complexity is in the Eye of the Beholder</i> Sian Gooding, Ekaterina Kochmar, Seid Muhie Yimam and Chris Biemann	4439
<i>Linguistic Complexity Loss in Text-Based Therapy</i> Jason Wei, Kelly Finn, Emma Templeton, Thalia Wheatley and Soroush Vosoughi	4450
<i>Ab Antiquo: Neural Proto-language Reconstruction</i> Carlo Meloni, Shauli Ravfogel and Yoav Goldberg	4460
<i>On Biasing Transformer Attention Towards Monotonicity</i> Annette Rios, Chantal Amrhein, Noëmi Aepli and Rico Sennrich	4474
<i>Extracting a Knowledge Base of Mechanisms from COVID-19 Papers</i> Tom Hope, Aida Amini, David Wadden, Madeleine van Zuylen, Sravanthi Parasa, Eric Horvitz, Daniel Weld, Roy Schwartz and Hannaneh Hajishirzi	4489
<i>Constrained Multi-Task Learning for Event Coreference Resolution</i> Jing Lu and Vincent Ng	4504

<i>Empirical Evaluation of Pre-trained Transformers for Human-Level NLP: The Role of Sample Size and Dimensionality</i>	
Adithya V Ganesan, Matthew Matero, Aravind Reddy Ravula, Huy Vu and H. Andrew Schwartz	4515
<i>Leveraging Deep Representations of Radiology Reports in Survival Analysis for Predicting Heart Failure Patient Mortality</i>	
Hyun Gi Lee, Evan Sholle, Ashley Beecy, Subhi Al’Aref and Yifan Peng	4533
<i>On the Use of Context for Predicting Citation Worthiness of Sentences in Scholarly Articles</i>	
Rakesh Gosangi, Ravneet Arora, Mohsen Gheisarieha, Debanjan Mahata and Haimin Zhang	4539
<i>Data and Model Distillation as a Solution for Domain-transferable Fact Verification</i>	
Mitch Paul Mithun, Sandeep Suntwal and Mihai Surdeanu	4546
<i>Adapting Coreference Resolution for Processing Violent Death Narratives</i>	
Ankith Uppunda, Susan Cochran, Jacob Foster, Alina Arseniev-Koehler, Vickie Mays and Kai-Wei Chang	4553
<i>Time-Stamped Language Model: Teaching Language Models to Understand The Flow of Events</i>	
Hossein Rajaby Faghihi and Parisa Kordjamshidi	4560
<i>If You Want to Go Far Go Together: Unsupervised Joint Candidate Evidence Retrieval for Multi-hop Question Answering</i>	
Vikas Yadav, Steven Bethard and Mihai Surdeanu	4571
<i>SPARTQA: A Textual Question Answering Benchmark for Spatial Reasoning</i>	
Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang Ning and Parisa Kordjamshidi	4582
<i>A Dataset of Information-Seeking Questions and Answers Anchored in Research Papers</i>	
Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith and Matt Gardner	4599
<i>Differentiable Open-Ended Commonsense Reasoning</i>	
Bill Yuchen Lin, Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Xiang Ren and William Cohen	4611
<i>Does Structure Matter? Encoding Documents for Machine Reading Comprehension</i>	
Hui Wan, Song Feng, Chulaka Gunasekara, Siva Sankalp Patel, Sachindra Joshi and Luis Lastras	4626
<i>Multi-Step Reasoning Over Unstructured Text with Beam Dense Retrieval</i>	
Chen Zhao, Chenyan Xiong, Jordan Boyd-Graber and Hal Daumé III	4635
<i>Scalable and Interpretable Semantic Change Detection</i>	
Syrielle Montariol, Matej Martinc and Lidia Pivovarova	4642
<i>Scalar Adjective Identification and Multilingual Ranking</i>	
Aina Garí Soler and Marianna Apidianaki	4653
<i>ESC: Redesigning WSD with Extractive Sense Comprehension</i>	
Edoardo Barba, Tommaso Pasini and Roberto Navigli	4661
<i>Recent advances in neural metaphor processing: A linguistic, cognitive and social perspective</i>	
Xiaoyu Tong, Ekaterina Shutova and Martha Lewis	4673

<i>Constructing Taxonomies from Pretrained Language Models</i> Catherine Chen, Kevin Lin and Dan Klein .....	4687
<i>Event Representation with Sequential, Semi-Supervised Discrete Variables</i> Mehdi Rezaee and Francis Ferraro .....	4701
<i>Seq2Emo: A Sequence to Multi-Label Emotion Classification Model</i> Chenyang Huang, Amine Trabelsi, Xuebin Qin, Nawshad Farruque, Lili Mou and Osmar Zaiane 4717	
<i>Knowledge Enhanced Masked Language Model for Stance Detection</i> Kornraphop Kawintiranon and Lisa Singh .....	4725
<i>Learning Paralinguistic Features from Audiobooks through Style Voice Conversion</i> Zakaria Aldeneh, Matthew Perez and Emily Mower Provost .....	4736
<i>Adapting BERT for Continual Learning of a Sequence of Aspect Sentiment Classification Tasks</i> Zixuan Ke, Hu Xu and Bing Liu .....	4746
<i>Adversarial Learning for Zero-Shot Stance Detection on Social Media</i> Emily Allaway, Malavika Srikanth and Kathleen McKeown .....	4756
<i>Efficiently Summarizing Text and Graph Encodings of Multi-Document Clusters</i> Ramakanth Pasunuru, Mengwen Liu, Mohit Bansal, Sujith Ravi and Markus Dreyer .....	4768
<i>Enriching Transformers with Structured Tensor-Product Representations for Abstractive Summarization</i> Yichen Jiang, Asli Celikyilmaz, Paul Smolensky, Paul Soulos, Sudha Rao, Hamid Palangi, Roland Fernandez, Caitlin Smith, Mohit Bansal and Jianfeng Gao .....	4780
<i>What's in a Summary? Laying the Groundwork for Advances in Hospital-Course Summarization</i> Griffin Adams, Emily Alsentzer, Mert Ketenci, Jason Zucker and Noémie Elhadad .....	4794
<i>Understanding Factuality in Abstractive Summarization with FRANK: A Benchmark for Factuality Metrics</i> Artidoro Pagnoni, Vidhisha Balachandran and Yulia Tsvetkov .....	4812
<i>GSum: A General Framework for Guided Neural Abstractive Summarization</i> Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang and Graham Neubig .....	4830
<i>What Will it Take to Fix Benchmarking in Natural Language Understanding?</i> Samuel R. Bowman and George Dahl .....	4843
<i>TuringAdvice: A Generative and Dynamic Evaluation of Language Use</i> Rowan Zellers, Ari Holtzman, Elizabeth Clark, Lianhui Qin, Ali Farhadi and Yejin Choi .....	4856
<i>Multitask Learning for Emotionally Analyzing Sexual Abuse Disclosures</i> Ramit Sawhney, Puneet Mathur, Taru Jain, Akash Kumar Gautam and Rajiv Ratn Shah .....	4881
<i>Self Promotion in US Congressional Tweets</i> Jun Wang, Kelly Cui and Bei Yu .....	4893
<i>Profiling of Intertextuality in Latin Literature Using Word Embeddings</i> Patrick J. Burns, James Brofos, Kyle Li, Pramit Chaudhuri and Joseph P. Dexter .....	4900
<i>Identifying inherent disagreement in natural language inference</i> Xinliang Frederick Zhang and Marie-Catherine de Marneffe .....	4908

<i>Modeling Human Mental States with an Entity-based Narrative Graph</i>	
I-Ta Lee, Maria Leonor Pacheco and Dan Goldwasser . . . . .	4916
<i>A Simple and Efficient Multi-Task Learning Approach for Conditioned Dialogue Generation</i>	
Yan Zeng and Jian-Yun Nie . . . . .	4927
<i>Hurdles to Progress in Long-form Question Answering</i>	
Kalpesh Krishna, Aurko Roy and Mohit Iyyer . . . . .	4940
<i>ENTRUST: Argument Reframing with Language Models and Entailment</i>	
Tuhin Chakrabarty, Christopher Hidey and Smaranda Muresan . . . . .	4958
<i>Paragraph-level Simplification of Medical Texts</i>	
Ashwin Devaraj, Iain Marshall, Byron Wallace and Junyi Jessy Li . . . . .	4972
<i>An Empirical Study on Neural Keyphrase Generation</i>	
Rui Meng, Xingdi Yuan, Tong Wang, Sanqiang Zhao, Adam Trischler and Daqing He . . . . .	4985
<i>Attention Head Masking for Inference Time Content Selection in Abstractive Summarization</i>	
Shuyang Cao and Lu Wang . . . . .	5008
<i>Factual Probing Is [MASK]: Learning vs. Learning to Recall</i>	
Zexuan Zhong, Dan Friedman and Danqi Chen . . . . .	5017
<i>Evaluating Saliency Methods for Neural Language Models</i>	
Shuoyang Ding and Philipp Koehn . . . . .	5034
<i>Contextualized Perturbation for Textual Adversarial Attack</i>	
Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun and Bill Dolan	
5053	
<i>DirectProbe: Studying Representations without Classifiers</i>	
Yichu Zhou and Vivek Srikumar . . . . .	5070
<i>Evaluating the Values of Sources in Transfer Learning</i>	
Md Rizwan Parvez and Kai-Wei Chang . . . . .	5084
<i>Too Much in Common: Shifting of Embeddings in Transformer Language Models and its Implications</i>	
Daniel Biś, Maksim Podkorytov and Xiuwen Liu . . . . .	5117
<i>On the Inductive Bias of Masked Language Modeling: From Statistical to Syntactic Dependencies</i>	
Tianyi Zhang and Tatsunori Hashimoto . . . . .	5131
<i>Limitations of Autoregressive Models and Their Alternatives</i>	
Chu-Cheng Lin, Aaron Jaech, Xin Li, Matthew R. Gormley and Jason Eisner . . . . .	5147
<i>On the Transformer Growth for Progressive BERT Training</i>	
Xiaotao Gu, Liyuan Liu, Hongkun Yu, Jing Li, Chen Chen and Jiawei Han . . . . .	5174
<i>Revisiting Simple Neural Probabilistic Language Models</i>	
Simeng Sun and Mohit Iyyer . . . . .	5181
<i>ReadTwice: Reading Very Large Documents with Memories</i>	
Yury Zemlyanskiy, Joshua Ainslie, Michiel de Jong, Philip Pham, Ilya Eckstein and Fei Sha .	5189

<i>SCRIPT: Self-Critic PreTraining of Transformers</i> Erik Nijkamp, Bo Pang, Ying Nian Wu and Caiming Xiong .....	5196
<i>Learning How to Ask: Querying LMs with Mixtures of Soft Prompts</i> Guanghui Qin and Jason Eisner .....	5203
<i>Nutri-bullets Hybrid: Consensual Multi-document Summarization</i> Darsh Shah, Lili Yu, Tao Lei and Regina Barzilay .....	5213
<i>AVA: an Automatic eValuation Approach for Question Answering Systems</i> Thuy Vu and Alessandro Moschitti .....	5223
<i>SpanPredict: Extraction of Predictive Document Spans with Neural Attention</i> Vivek Subramanian, Matthew Engelhard, Sam Berchuck, Liqun Chen, Ricardo Henao and Lawrence Carin .....	5234
<i>Text Editing by Command</i> Felix Faltings, Michel Galley, Gerold Hintz, Chris Brockett, Chris Quirk, Jianfeng Gao and Bill Dolan .....	5259
<i>A Deep Metric Learning Approach to Account Linking</i> Aleem Khan, Elizabeth Fleming, Noah Schofield, Marcus Bishop and Nicholas Andrews .....	5275
<i>Improving Factual Completeness and Consistency of Image-to-Text Radiology Report Generation</i> Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz and Dan Jurafsky .....	5288
<i>Multimodal End-to-End Sparse Model for Emotion Recognition</i> Wenliang Dai, Samuel Cahyawijaya, Zihan Liu and Pascale Fung .....	5305
<i>MIMOQA: Multimodal Input Multimodal Output Question Answering</i> Hrituraj Singh, Anshul Nasery, Denil Mehta, Aishwarya Agarwal, Jatin Lamba and Balaji Vasanth Srinivasan .....	5317
<i>OCID-Ref: A 3D Robotic Dataset With Embodied Language For Clutter Scene Grounding</i> Ke-Jyun Wang, Yun-Hsuan Liu, Hung-Ting Su, Jen-Wei Wang, Yu-Siang Wang, Winston Hsu and Wen-Chin Chen .....	5333
<i>Unsupervised Vision-and-Language Pre-training Without Parallel Images and Captions</i> Lunian Harold Li, Haoxuan You, Zhecan Wang, Alireza Zareian, Shih-Fu Chang and Kai-Wei Chang .....	5339
<i>Multitasking Inhibits Semantic Drift</i> Athul Paul Jacob, Mike Lewis and Jacob Andreas .....	5351
<i>Probing Contextual Language Models for Common Ground with Visual Representations</i> Gabriel Ilharco, Rowan Zellers, Ali Farhadi and Hannaneh Hajishirzi .....	5367
<i>BBAEG: Towards BERT-based Biomedical Adversarial Example Generation for Text Classification</i> Ishani Mondal .....	5378
<i>Targeted Adversarial Training for Natural Language Understanding</i> Lis Pereira, Xiaodong Liu, Hao Cheng, Hoifung Poon, Jianfeng Gao and Ichiro Kobayashi ..	5385
<i>Latent-Optimized Adversarial Neural Transfer for Sarcasm Detection</i> Xu Guo, Boyang Li, Han Yu and Chunyan Miao .....	5394

<i>Self-training Improves Pre-training for Natural Language Understanding</i>	
Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Veselin Stoyanov and Alexis Conneau .....	5408
<i>Supporting Clustering with Contrastive Learning</i>	
Dejiao Zhang, Feng Nan, Xiaokai Wei, Shang-Wen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew O. Arnold and Bing Xiang .....	5419
<i>TITA: A Two-stage Interaction and Topic-Aware Text Matching Model</i>	
Xingwu Sun, Yanling Cui, Hongyin Tang, Qiuyu Zhu, Fuzheng Zhang and Beihong Jin .....	5431
<i>Neural Quality Estimation with Multiple Hypotheses for Grammatical Error Correction</i>	
Zhenghao Liu, Xiaoyuan Yi, Maosong Sun, Liner Yang and Tat-Seng Chua .....	5441
<i>Neural Network Surgery: Injecting Data Patterns into Pre-trained Models with Minimal Instance-wise Side Effects</i>	
Zhiyuan Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He .....	5453
<i>Discrete Argument Representation Learning for Interactive Argument Pair Identification</i>	
Lu Ji, Zhongyu Wei, Jing Li, Qi Zhang and Xuanjing Huang .....	5467
<i>On Unifying Misinformation Detection</i>	
Nayeon Lee, Belinda Z. Li, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih and Madian Khabsa	5479
<i>Frustratingly Easy Edit-based Linguistic Steganography with a Masked Language Model</i>	
Honai Ueoka, Yugo Murawaki and Sadao Kurohashi .....	5486
<i>Few-Shot Text Classification with Triplet Networks, Data Augmentation, and Curriculum Learning</i>	
Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng and Shiqi Xu .....	5493
<i>Do RNN States Encode Abstract Phonological Alternations?</i>	
Miikka Silfverberg, Francis Tyers, Garrett Nicolai and Mans Hulden .....	5501
<i>Pre-training with Meta Learning for Chinese Word Segmentation</i>	
Zhen Ke, Liang Shi, Songtao Sun, Erli Meng, Bin Wang and Xipeng Qiu .....	5514
<i>Decompose, Fuse and Generate: A Formation-Informed Method for Chinese Definition Generation</i>	
Hua Zheng, Damai Dai, Lei Li, Tianyu Liu, Zhifang Sui, Baobao Chang and Yang Liu .....	5524
<i>User-Generated Text Corpus for Evaluating Japanese Morphological Analysis and Lexical Normalization</i>	
Shohei Higashiyama, Masao Utiyama, Taro Watanabe and Eiichiro Sumita .....	5532
<i>GPT Perdetry Test: Generating new meanings for new words</i>	
Nikolay Malkin, Sameera Lanka, Pranav Goel, Sudha Rao and Nebojsa Jojic .....	5542
<i>Universal Semantic Tagging for English and Mandarin Chinese</i>	
Wenxi Li, Yiyang Hou, Yajie Ye, Li Liang and Weiwei Sun .....	5554
<i>ShadowGNN: Graph Projection Neural Network for Text-to-SQL Parser</i>	
Zhi Chen, Lu Chen, Yanbin Zhao, Ruisheng Cao, Zihan Xu, Su Zhu and Kai Yu .....	5567
<i>Contextualized and Generalized Sentence Representations by Contrastive Self-Supervised Learning: A Case Study on Discourse Relation Analysis</i>	
Hirokazu Kiyomaru and Sadao Kurohashi .....	5578



<i>AMR Parsing with Action-Pointer Transformer</i>	
Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo and Radu Florian . . . . .	5585
<i>NL-EDIT: Correcting Semantic Parse Errors through Natural Language Interaction</i>	
Ahmed Elgohary, Christopher Meek, Matthew Richardson, Adam Fourney, Gonzalo Ramos and Ahmed Hassan Awadallah . . . . .	5599
<i>Unsupervised Concept Representation Learning for Length-Varying Text Similarity</i>	
Xuchao Zhang, Bo Zong, Wei Cheng, Jingchao Ni, Yanchi Liu and Haifeng Chen . . . . .	5611
<i>Augmenting Knowledge-grounded Conversations with Sequential Knowledge Transition</i>	
Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao and Yanyan Lan . .	5621
<i>Adversarial Self-Supervised Learning for Out-of-Domain Detection</i>	
Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu and Weiran Xu . . . . .	5631
<i>Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue State Tracking</i>	
Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho and Rajen Subba . . . . .	5640
<i>Hierarchical Transformer for Task Oriented Dialog Systems</i>	
Bishal Santra, Potnuru Anusha and Pawan Goyal . . . . .	5649
<i>Measuring the ‘I don’t know’ Problem through the Lens of Gricean Quantity</i>	
Huda Khayrallah and João Sedoc . . . . .	5659
<i>RTFE: A Recursive Temporal Fact Embedding Framework for Temporal Knowledge Graph Completion</i>	
Youri Xu, Haihong E, Meina Song, wenyu song, Xiaodong Lv, wang haotian and yang jinrui .	5671
<i>Open Hierarchical Relation Extraction</i>	
Kai Zhang, Yuan Yao, Ruobing Xie, Xu Han, Zhiyuan Liu, Fen Lin, Leyu Lin and Maosong Sun	5682
<i>Jointly Extracting Explicit and Implicit Relational Triples with Reasoning Pattern Enhanced Binary Pointer Network</i>	
Yubo Chen, Yunqi Zhang, Changran Hu and Yongfeng Huang . . . . .	5694
<i>Multi-Grained Knowledge Distillation for Named Entity Recognition</i>	
Xuan Zhou, Xiao Zhang, Chenyang Tao, Junya Chen, Bing Xu, Wei Wang and Jing Xiao . . . .	5704
<i>SGG: Learning to Select, Guide, and Generate for Keyphrase Generation</i>	
Jing Zhao, Junwei Bao, Yifan Wang, Youzheng Wu, Xiaodong He and Bowen Zhou . . . . .	5717
<i>Towards Sentiment and Emotion aided Multi-modal Speech Act Classification in Twitter</i>	
Tulika Saha, Apoorva Upadhyaya, Sriparna Saha and Pushpak Bhattacharyya . . . . .	5727
<i>Generative Imagination Elevates Machine Translation</i>	
Quanyu Long, Mingxuan Wang and Lei Li . . . . .	5738
<i>Non-Autoregressive Translation by Learning Target Categorical Codes</i>	
Yu Bao, Shujian Huang, Tong Xiao, Dongqi Wang, Xinyu Dai and Jiajun CHEN . . . . .	5749
<i>Training Data Augmentation for Code-Mixed Translation</i>	
Abhirut Gupta, Aditya Vavre and Sunita Sarawagi . . . . .	5760

<i>Rethinking Perturbations in Encoder-Decoders for Fast Training</i> Sho Takase and Shun Kiyono .....	5767
<i>Context-aware Decoder for Neural Machine Translation using a Target-side Document-Level Language Model</i> Amane Sugiyama and Naoki Yoshinaga.....	5781
<i>Machine Translated Text Detection Through Text Similarity with Round-Trip Translation</i> Hoang-Quoc Nguyen-Son, Tran Thao, Seira Hidano, Ishita Gupta and Shinsaku Kiyomoto ..	5792
<i>TR-BERT: Dynamic Token Reduction for Accelerating BERT Inference</i> Deming Ye, Yankai Lin, Yufei Huang and Maosong Sun.....	5798
<i>Breadth First Reasoning Graph for Multi-hop Question Answering</i> Yongjie Huang and Meng Yang .....	5810
<i>Improving Zero-Shot Cross-lingual Transfer for Multilingual Question Answering over Knowledge Graph</i> Yucheng Zhou, Xiubo Geng, Tao Shen, Wenqiang Zhang and Daxin Jiang.....	5822
<i>RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering</i> Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu and Haifeng Wang .....	5835
<i>DAGN: Discourse-Aware Graph Network for Logical Reasoning</i> Yinya Huang, Meng Fang, Yu Cao, Liwei Wang and Xiaodan Liang .....	5848
<i>Designing a Minimal Retrieve-and-Read System for Open-Domain Question Answering</i> Sohee Yang and Minjoon Seo .....	5856
<i>Unsupervised Multi-hop Question Answering by Question Generation</i> Liangming Pan, Wenhua Chen, Wenhan Xiong, Min-Yen Kan and William Yang Wang.....	5866
<i>Sliding Selector Network with Dynamic Memory for Extractive Summarization of Long Documents</i> Peng Cui and Le Hu .....	5881
<i>AdaptSum: Towards Low-Resource Domain Adaptation for Abstractive Summarization</i> Tiezheng Yu, Zihan Liu and Pascale Fung .....	5892
<i>QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization</i> Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu and Dragomir Radev .....	5905
<i>MM-AVS: A Full-Scale Dataset for Multi-modal Summarization</i> Xiyan Fu, Jun Wang and Zhenglu Yang .....	5922
<i>MediaSum: A Large-scale Media Interview Dataset for Dialogue Summarization</i> Chenguang Zhu, Yang Liu, Jie Mei and Michael Zeng .....	5927
<i>Improving Faithfulness in Abstractive Summarization with Contrast Candidate Generation and Selection</i> Sihao Chen, Fan Zhang, Kazuo Sone and Dan Roth .....	5935
<i>Inference Time Style Control for Summarization</i> Shuyang Cao and Lu Wang .....	5942

# Conference Program

Mon 07 Jun 2021 (all times PDT, UTC-7)

08:00–09:00 Keynote

09:00–10:20 1A: Information Extraction

*Knowledge Router: Learning Disentangled Representations for Knowledge Graphs*  
Shuai Zhang, Xi Rao, Yi Tay and Ce Zhang

*Distantly Supervised Relation Extraction with Sentence Reconstruction and Knowledge Base Priors*  
Fenia Christopoulou, Makoto Miwa and Sophia Ananiadou

*Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks*  
Minh Van Nguyen, Viet Lai and Thien Huu Nguyen

*Abstract Meaning Representation Guided Graph Encoding and Decoding for Joint Information Extraction*  
Zixuan Zhang and Heng Ji

*A Frustratingly Easy Approach for Entity and Relation Extraction*  
Zexuan Zhong and Danqi Chen

*Event Time Extraction and Propagation via Graph Attention Networks*  
Haoyang Wen, Yanru Qu, Heng Ji, Qiang Ning, Jiawei Han, Avi Sil, Hanghang Tong and Dan Roth

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

09:00–10:20 1B: Interpretability and Analysis of Models for NLP

*Probing Word Translations in the Transformer and Trading Decoder for Encoder Layers*

Hongfei Xu, Josef van Genabith, Qihui Liu and Deyi Xiong

*Mediators in Determining what Processing BERT Performs First*

Aviv Slobodkin, Leshem Choshen and Omri Abend

*Automatic Generation of Contrast Sets from Scene Graphs: Probing the Compositional Consistency of GQA*

Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz and Michael Elhadad

*Multilingual Language Models Predict Human Reading Behavior*

Nora Hollenstein, Federico Pirovano, Ce Zhang, Lena Jäger and Lisa Beinborn

*Do Syntactic Probes Probe Syntax? Experiments with Jabberwocky Probing*

Rowan Hall Maudslay and Ryan Cotterell

*A Non-Linear Structural Probe*

Jennifer C. White, Tiago Pimentel, Naomi Saphra and Ryan Cotterell

*Concealed Data Poisoning Attacks on NLP Models*

Eric Wallace, Tony Zhao, Shi Feng and Sameer Singh

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

09:00–10:20 1C: Machine Translation

*Backtranslation Feedback Improves User Confidence in MT, Not Quality*

Vilém Zouhar, Michal Novák, Matuš Žilínek, Ondřej Bojar, Mateo Obregón, Robin L. Hill, Frédéric Blain, Marina Fomicheva, Lucia Specia and Lisa Yankovskaya

*Data Filtering using Cross-Lingual Word Embeddings*

Christian Herold, Jan Rosendahl, Joris Vanvinckenroye and Hermann Ney

*Improving the Lexical Ability of Pretrained Language Models for Unsupervised Neural Machine Translation*

Alexandra Chronopoulou, Dario Stojanovski and Alexander Fraser

*Neural Machine Translation without Embeddings*

Uri Shaham and Omer Levy

*Counterfactual Data Augmentation for Neural Machine Translation*

Qi Liu, Matt Kusner and Phil Blunsom

*Cultural and Geographical Influences on Image Translatability of Words across Languages*

Nikzad Khani, Isidora Tourni, Mohammad Sadegh Rasooli, Chris Callison-Burch and Derry Tanti Wijaya

*Multilingual BERT Post-Pretraining Alignment*

Lin Pan, Chung-Wei Hang, Haode Qi, Abhishek Shah, Saloni Potdar and Mo Yu

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

09:00–10:20 1D: NLP Applications

*A Million Tweets Are Worth a Few Points: Tuning Transformers for Customer Service Tasks*

Amir Hadifar, Sofie Labat, Veronique Hoste, Chris Develder and Thomas De-meester

*Paragraph-level Rationale Extraction through Regularization: A case study on European Court of Human Rights Cases*

Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos and Prodromos Malakasiotis

*Answering Product-Questions by Utilizing Questions from Other Contextually Similar Products*

Ohad Rozen, David Carmel, Avihai Mejer, Vitaly Mirkis and Yftah Ziser

*EnSidNet: Enhanced Hybrid Siamese-Deep Network for grouping clinical trials into drug-development pathways*

Lucia Pagani

*DATE: Detecting Anomalies in Text via Self-Supervision of Transformers*

Andrei Manolache, Florin Brad and Elena Burceanu

*A Simple Approach for Handling Out-of-Vocabulary Identifiers in Deep Learning for Source Code*

Nadezhda Chirkova and Sergey Troshin

*Fast and Scalable Dialogue State Tracking with Explicit Modular Decomposition*

Dingmin Wang, Chenghua Lin, Qi Liu and Kam-Fai Wong

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**09:00–10:20 1E: Sentence-level Semantics and Textual Inference**

*Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks*

Nandan Thakur, Nils Reimers, Johannes Daxenberger and Iryna Gurevych

*SmBoP: Semi-autoregressive Bottom-up Semantic Parsing*

Ohad Rubín and Jonathan Berant

*SGL: Speaking the Graph Languages of Semantic Parsing via Multilingual Translation*

Luigi Procopio, Rocco Tripodi and Roberto Navigli

*Unifying Cross-Lingual Semantic Role Labeling with Heterogeneous Linguistic Resources*

Simone Conia, Andrea Bacciu and Roberto Navigli

*Fool Me Twice: Entailment from Wikipedia Gamification*

Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger and Jordan Boyd-Graber

*Meta-Learning for Domain Generalization in Semantic Parsing*

Bailin Wang, Mirella Lapata and Ivan Titov

**10:20–11:40 2A: Language Generation**

*Aspect-Controlled Neural Argument Generation*

Benjamin Schiller, Johannes Daxenberger and Iryna Gurevych

*Text Generation from Discourse Representation Structures*

Jiangming Liu, Shay B. Cohen and Mirella Lapata

*APo-VAE: Text Generation in Hyperbolic Space*

Shuyang Dai, Zhe Gan, Yu Cheng, Chenyang Tao, Lawrence Carin and Jingjing Liu

*DART: Open-Domain Structured Data Record to Text Generation*

Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chian-chun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiyaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher and Nazneen Fatema Rajani

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*[TACL14] An Error Analysis Framework for Shallow Surface Realisation*  
Shimorina, Anastasia, Parmentier, Yannick, Gardent, Claire

**10:20–11:40 2B: Multilinguality**

*When Being Unseen from mBERT is just the Beginning: Handling New Languages With Multilingual Language Models*

Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot and Djamel Seddah

*Multi-Adversarial Learning for Cross-Lingual Word Embeddings*

Haozhou Wang, James Henderson and Paola Merlo

*Multi-view Subword Regularization*

Xinyi Wang, Sebastian Ruder and Graham Neubig

*mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer*

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua and Colin Raffel

*MetaXL: Meta Representation Transformation for Low-resource Cross-lingual Learning*

Mengzhou Xia, Guoqing Zheng, Subhabrata Mukherjee, Milad Shokouhi, Graham Neubig and Ahmed Hassan Awadallah

*[TACL1] Parameter Space Factorization for Zero-Shot Learning across Tasks and Languages*

Edoardo M. Ponti, Ivan Vulić, Ryan Cotterell, Marinela Parovic, Roi Reichart, Anna Korhonen

**10:20–11:40 2C: Question Answering**

*Open Domain Question Answering over Tables via Dense Retrieval*

Jonathan Herzig, Thomas Müller, Syrine Krichene and Julian Eisenschlos

*Open-Domain Question Answering Goes Conversational via Question Rewriting*

Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, Shayne Longpre, Stephen Pulman and Srinivas Chappidi

*QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering*

Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang and Jure Leskovec



**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*XOR QA: Cross-lingual Open-Retrieval Question Answering*

Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi and Hannaneh Hajishirzi

*SPARTA: Efficient Open-Domain Question Answering via Sparse Transformer Matching Retrieval*

Tiancheng Zhao, Xiaopeng Lu and Kyusong Lee

*[TACL12] Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies*

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, Jonathan Berant

**10:20–11:40 2D: Special Theme: New Challenges in NLP**

*Implicitly Abusive Language – What does it actually look like and why are we not getting there?*

Michael Wiegand, Josef Ruppenhofer and Elisabeth Eder

*The Importance of Modeling Social Factors of Language: Theory and Practice*

Dirk Hovy and Diyi Yang

*On learning and representing social meaning in NLP: a sociolinguistic perspective*

Dong Nguyen, Laura Rosseel and Jack Grieve

*Preregistering NLP research*

Emiel van Miltenburg, Chris van der Lee and Emiel Krahmer

*Get Your Vitamin C! Robust Fact Verification with Contrastive Evidence*

Tal Schuster, Adam Fisch and Regina Barzilay

*Representing Numbers in NLP: a Survey and a Vision*

Avijit Thawani, Jay Pujara, Filip Ilievski and Pedro Szekely

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**10:20–11:40 2E: Summarization**

*Extending Multi-Document Summarization Evaluation to the Interactive Setting*

Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer and Ido Dagan

*Identifying Helpful Sentences in Product Reviews*

Iftah Gamzu, Hila Gonen, Gilad Kutiel, Ran Levy and Eugene Agichtein

*Noisy Self-Knowledge Distillation for Text Summarization*

Yang Liu, Sheng Shen and Mirella Lapata

*Improving Zero and Few-Shot Abstractive Summarization with Intermediate Fine-tuning and Data Augmentation*

Alexander Fabbri, Simeng Han, Haoyuan Li, Haoran Li, Marjan Ghazvininejad, Shafiq Joty, Dragomir Radev and Yashar Mehdad

*Enhancing Factual Consistency of Abstractive Summarization*

Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang and Meng Jiang

*[TACL6] Extractive Opinion Summarization in Quantized Transformer Spaces*

Stefanos Angelidis, Reinald Kim Amplayo, Yoshihiko Suhara, Xiaolan Wang, Mirella Lapata

**11:40–13:00 3A: Dialogue and Interactive Systems**

*Few-shot Intent Classification and Slot Filling with Retrieved Examples*

Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat and Qi Li

*"Nice Try, Kiddo": Investigating Ad Hominems in Dialogue Responses*

Emily Sheng, Kai-Wei Chang, Prem Natarajan and Nanyun Peng

*Human-like informative conversations: Better acknowledgements using conditional mutual information*

Ashwin Paranjape and Christopher Manning

*A Comparative Study on Schema-Guided Dialogue State Tracking*

Jie Cao and Yi Zhang

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*Spoken Language Understanding for Task-oriented Dialogue Systems with Augmented Memory Networks*

Jie Wu, Ian Harris and Hongzhi Zhao

*How to Motivate Your Dragon: Teaching Goal-Driven Agents to Speak and Act in Fantasy Worlds*

Prithviraj Ammanabrolu, Jack Urbanek, Margaret Li, Arthur Szlam, Tim Rocktäschel and Jason Weston

**11:40–13:00 3B: Information Extraction**

*Linking Entities to Unseen Knowledge Bases with Arbitrary Schemas*

Yogarshi Vyas and Miguel Ballesteros

*Self-Training with Weak Supervision*

Giannis Karamanolakis, Subhabrata Mukherjee, Guoqing Zheng and Ahmed Hassan Awadallah

*Neural Language Modeling for Contextualized Temporal Graph Generation*

Aman Madaan and Yiming Yang

*Probabilistic Box Embeddings for Uncertain Knowledge Graph Reasoning*

Xuelu Chen, Michael Boratko, Muhao Chen, Shib Sankar Dasgupta, Xiang Lorraine Li and Andrew McCallum

*Document-Level Event Argument Extraction by Conditional Generation*

Sha Li, Heng Ji and Jiawei Han

*Template Filling with Generative Transformers*

Xinya Du, Alexander Rush and Claire Cardie

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

11:40–13:00 **3C: Interpretability and Analysis of Models for NLP**

*Towards Interpreting and Mitigating Shortcut Learning Behavior of NLU models*

Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun and Xia Hu

*On Attention Redundancy: A Comprehensive Study*

Yuchen Bian, Jiaji Huang, Xingyu Cai, Jiahong Yuan and Kenneth Church

*Does BERT Pretrained on Clinical Notes Reveal Sensitive Data?*

Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg and Byron Wallace

*Low-Complexity Probing via Finding Subnetworks*

Victor Sanh and Alexander Rush

*An Empirical Comparison of Instance Attribution Methods for NLP*

Pouya Pezeshkpour, Sarthak Jain, Byron Wallace and Sameer Singh

*Generalization in Instruction Following Systems*

Soham Dan, Michael Zhou and Dan Roth

*[CL2] Interpretability Analysis for Named Entity Recognition to Understand System Predictions and How They Can Improve*

Oshin Agarwal, Yinfei Yang, Byron C. Wallace, Ani Nenkova

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**11:40–13:00 3D: Language Grounding to Vision, Robotics and Beyond**

*LightningDOT: Pre-training Visual-Semantic Embeddings for Real-Time Image-Text Retrieval*

Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang and Jingjing Liu

*Measuring Social Biases in Grounded Vision and Language Embeddings*

Candace Ross, Boris Katz and Andrei Barbu

*MTAG: Modal-Temporal Attention Graph for Unaligned Human Multimodal Language Sequences*

Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria and Louis-Philippe Morency

*Grounding Open-Domain Instructions to Automate Web Support Tasks*

Nancy Xu, Sam Masling, Michael Du, Giovanni Campagna, Larry Heck, James Landay and Monica Lam

*Modular Networks for Compositional Instruction Following*

Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein and trevor darrell

*Improving Cross-Modal Alignment in Vision Language Navigation via Syntactic Information*

Jialu Li, Hao Tan and Mohit Bansal

**11:40–13:00 3E: Machine Learning for NLP: Classification and Structured Prediction Models**

*Improving Pretrained Models for Zero-shot Multi-label Text Classification through Reinforced Label Hierarchy Reasoning*

Hui Liu, Danqing Zhang, Bing Yin and Xiaodan Zhu

*Fine-Tuning Pre-trained Language Model with Weak Supervision: A Contrastive-Regularized Self-Training Approach*

Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao and Chao Zhang

*Posterior Differential Regularization with  $f$ -divergence for Improving Model Robustness*

Hao Cheng, Xiaodong Liu, Lis Pereira, Yaoliang Yu and Jianfeng Gao

*Understanding Hard Negatives in Noise Contrastive Estimation*

Wenzheng Zhang and Karl Stratos

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*Certified Robustness to Word Substitution Attack with Differential Privacy*

Wenjie Wang, Pengfei Tang, Jian Lou and Li Xiong

*DReCa: A General Task Augmentation Strategy for Few-Shot Natural Language Inference*

Shikhar Murty, Tatsunori Hashimoto and Christopher Manning

**16:00–17:00 Keynote**

**17:00–18:20 4A: Machine Translation**

*Harnessing Multilinguality in Unsupervised Machine Translation for Rare Languages*

Xavier Garcia, Aditya Siddhant, Orhan Firat and Ankur Parikh

*Macro-Average: Rare Types Are Important Too*

Thamme Gowda, Weiqiu You, Constantine Lignos and Jonathan May

*Assessing Reference-Free Peer Evaluation for Machine Translation*

Sweta Agrawal, George Foster, Markus Freitag and Colin Cherry

*The Curious Case of Hallucinations in Neural Machine Translation*

Vikas Raunak, Arul Menezes and Marcin Junczys-Dowmunt

*Towards Continual Learning for Multilingual Machine Translation via Vocabulary Substitution*

Xavier Garcia, Noah Constant, Ankur Parikh and Orhan Firat

*Towards Modeling the Style of Translators in Neural Machine Translation*

Yue Wang, Cuong Hoang and Marcello Federico

*[TACL4] Unsupervised Bitext Mining and Translation via Self-trained Contextual Embeddings*

Phillip Keung, Julian Salazar, Yichao Lu, Noah A. Smith

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

17:00–18:20 4B: Question Answering

*Self-Supervised Test-Time Learning for Reading Comprehension*

Pratyay Banerjee, Tejas Gokhale and Chitta Baral

*Capturing Row and Column Semantics in Transformer Based Question Answering over Tables*

Michael Glass, Mustafa Canim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avi Sil, Feifei Pan, Samarth Bharadwaj and Nicolas Rodolfo Fauceglia

*Explainable Multi-hop Verbal Reasoning Through Internal Monologue*

Zhengzhong Liang, Steven Bethard and Mihai Surdeanu

*Robust Question Answering Through Sub-part Alignment*

Jifan Chen and Greg Durrett

*Text Modular Networks: Learning to Decompose Tasks in the Language of Existing Models*

Tushar Khot, Daniel Khashabi, Kyle Richardson, Peter Clark and Ashish Sabharwal

*RECONSIDER: Improved Re-Ranking using Span-Focused Cross-Attention for Open Domain Question Answering*

Srinivasan Iyer, Sewon Min, Yashar Mehdad and Wen-tau Yih

*On the Transferability of Minimal Prediction Preserving Inputs in Question Answering*

Shayne Longpre, Yi Lu and Chris DuBois

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

17:00–18:20 4C: Sentence-level Semantics and Textual Inference

*Understanding by Understanding Not: Modeling Negation in Language Models*

Arian Hosseini, Siva Reddy, Dzmitry Bahdanau, R Devon Hjelm, Alessandro Sordani and Aaron Courville

*DuoRAT: Towards Simpler Text-to-SQL Models*

Torsten Scholak, Raymond Li, Dzmitry Bahdanau, Harm de Vries and Chris Pal

*Looking Beyond Sentence-Level Natural Language Inference for Question Answering and Text Summarization*

Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Lorraine Li, Pavan Kapanipathi and Kartik Talamadupula

*Structure-Grounded Pretraining for Text-to-SQL*

Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun and Matthew Richardson

*Incremental Few-shot Text Classification with Multi-round New Classes: Formulation, Dataset and System*

Congying Xia, Wenpeng Yin, Yihao Feng and Philip Yu

*Temporal Reasoning on Implicit Events from Distant Supervision*

Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal and Dan Roth

*Disentangling Semantics and Syntax in Sentence Embeddings with Pre-trained Language Models*

James Y. Huang, Kuan-Hao Huang and Kai-Wei Chang



Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**17:00–18:20 4D: Summarization**

*Structure-Aware Abstractive Conversation Summarization via Discourse and Action Graphs*

Jiaao Chen and Diyi Yang

*A New Approach to Overgenerating and Scoring Abstractive Summaries*

Kaiqiang Song, Bingqing Wang, Zhe Feng and Fei Liu

*D2S: Document-to-Slide Generation Via Query-Based Text Summarization*

Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang and Nancy X. R. Wang

*Efficient Attentions for Long Document Summarization*

Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji and Lu Wang

*RefSum: Refactoring Neural Summarization*

Yixin Liu, Zi-Yi Dou and Pengfei Liu

*Annotating and Modeling Fine-grained Factuality in Summarization*

Tanya Goyal and Greg Durrett

**17:00–18:20 4E: Syntax: Tagging, Chunking, and Parsing**

*Larger-Context Tagging: When and Why Does It Work?*

Jinlan Fu, Liangjing Feng, Qi Zhang, Xuanjing Huang and Pengfei Liu

*Neural Sequence Segmentation as Determining the Leftmost Segments*

Yangming Li, Lema Liu and Kaisheng Yao

*PCFGs Can Do Better: Inducing Probabilistic Context-Free Grammars with Many Symbols*

Songlin Yang, Yanpeng Zhao and Kewei Tu

*GEMNET: Effective Gated Gazetteer Representations for Recognizing Complex Entities in Low-context Input*

Tao Meng, Anjie Fang, Oleg Rokhlenko and Shervin Malmasi

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*Video-aided Unsupervised Grammar Induction*

Songyang Zhang, Linfeng Song, Lifeng Jin, Kun Xu, Dong Yu and Jiebo Luo

*[CL3] Universal Dependencies*

Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, Daniel Zeman

**18:20–19:40 5A: Dialogue and Interactive Systems**

*Generating Negative Samples by Manipulating Golden Responses for Unsupervised Learning of a Response Evaluation Model*

ChaeHun Park, Eugene Jang, Wonsuk Yang and Jong Park

*How Robust are Fact Checking Systems on Colloquial Claims?*

Byeongchang Kim, Hyunwoo Kim, Seokhee Hong and Gunhee Kim

*Fine-grained Post-training for Improving Retrieval-based Dialogue Systems*

Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko and Jungyun Seo

*Put Chatbot into Its Interlocutor's Shoes: New Framework to Learn Chatbot Responding with Intention*

Hsuan Su, Jiun-Hao Jhan, Fan-yun Sun, Saurav Sahay and Hung-yi Lee

*Adding Chit-Chat to Enhance Task-Oriented Dialogues*

Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho and Claire Cardie

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

18:20–19:40 5B: Discourse and Pragmatics

*Incorporating Syntax and Semantics in Coreference Resolution with Heterogeneous Graph Attention Network*

Fan Jiang and Trevor Cohn

*Context Tracking Network: Graph-based Context Modeling for Implicit Discourse Relation Recognition*

Yingxue Zhang, Fandong Meng, Peng Li, Ping Jian and Jie Zhou

*Improving Neural RST Parsing Model with Silver Agreement Subtrees*

Naoki Kobayashi, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura and Masaaki Nagata

*RST Parsing from Scratch*

Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty and Xiaoli Li

*Did they answer? Subjective acts and intents in conversational discourse*

Elisa Ferracane, Greg Durrett, Junyi Jessy Li and Katrin Erk

*Evaluating the Impact of a Hierarchical Discourse Representation on Entity Coreference Resolution Performance*

Sopan Khosla, James Fiacco and Carolyn Rosé

*Bridging Resolution: Making Sense of the State of the Art*

Hideo Kobayashi and Vincent Ng

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**18:20–19:40 5C: Machine Learning for NLP: Language Modeling and Sequence to Sequence Models**

*Explicitly Modeling Syntax in Language Models with Incremental Parsing and a Dynamic Oracle*

Yikang Shen, Shawn Tan, Alessandro Sordani, Siva Reddy and Aaron Courville

*Revisiting the Weaknesses of Reinforcement Learning for Neural Machine Translation*

Samuel Kiegeand and Julia Kreutzer

*Learning to Organize a Bag of Words into Sentences with Neural Networks: An Empirical Study*

Chongyang Tao, Shen Gao, Juntao Li, Yansong Feng, Dongyan Zhao and Rui Yan

*Mask Attention Networks: Rethinking and Strengthen Transformer*

Zhihao Fan, Yeyun Gong, Dayiheng Liu, Zhongyu Wei, Siyuan Wang, Jian Jiao, Nan Duan, Ruofei Zhang and Xuanjing Huang

*ERNIE-Gram: Pre-Training with Explicitly N-Gram Masked Language Modeling for Natural Language Understanding*

Dongling Xiao, Yu-Kun Li, Han Zhang, Yu Sun, Hao Tian, Hua Wu and Haifeng Wang

*Lattice-BERT: Leveraging Multi-Granularity Representations in Chinese Pre-trained Language Models*

Yuxuan Lai, Yijia Liu, Yansong Feng, Songfang Huang and Dongyan Zhao

**18:20–19:40 5D: Lexical Semantics**

*Modeling Event Plausibility with Consistent Conceptual Abstraction*

Ian Porada, Kaheer Suleman, Adam Trischler and Jackie Chi Kit Cheung

*UmlsBERT: Clinical Domain Knowledge Augmentation of Contextual Embeddings Using the Unified Medical Language System Metathesaurus*

George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen and Alexander Wong

*Field Embedding: A Unified Grain-Based Framework for Word Representation*

Junjie Luo, Xi Chen, Jichao Sun, Yuejia Xiang, Ningyu Zhang and Xiang Wan

*MelBERT: Metaphor Detection via Contextualized Late Interaction using Metaphorical Identification Theories*

Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee and Jongwuk Lee

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*Non-Parametric Few-Shot Learning for Word Sense Disambiguation*

Howard Chen, Mengzhou Xia and Danqi Chen

**18:20–19:40 5E: Sentiment Analysis and Stylistic Analysis**

*Why Do Document-Level Polarity Classifiers Fail?*

Karen Martins, Pedro O.S Vaz-de-Melo and Rodrygo Santos

*A Unified Span-Based Approach for Opinion Mining with Syntactic Constituents*

Qingrong Xia, Bo Zhang, Rui Wang, Zhenghua Li, Yue Zhang, Fei Huang, Luo Si and Min Zhang

*Target-specified Sequence Labeling with Multi-head Self-attention for Target-oriented Opinion Words Extraction*

Yuhao Feng, Yanghui Rao, Yuyao Tang, Ninghua Wang and He Liu

*Does syntax matter? A strong baseline for Aspect-based Sentiment Analysis with RoBERTa*

Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu and Xipeng Qiu

*Domain Divergences: A Survey and Empirical Analysis*

Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan and Roger Zimmermann

*Target-Aware Data Augmentation for Stance Detection*

Yingjie Li and Cornelia Caragea

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**19:40–21:00 6A: Speech**

*End-to-end ASR to jointly predict transcriptions and linguistic annotations*

Motoi Omachi, Yuya Fujita, Shinji Watanabe and Matthew Wiesner

*Source and Target Bidirectional Knowledge Distillation for End-to-end Speech Translation*

Hirofumi Inaguma, Tatsuya Kawahara and Shinji Watanabe

*Searchable Hidden Intermediates for End-to-End Models of Decomposable Sequence Tasks*

Siddharth Dalmia, Brian Yan, Vikas Raunak, Florian Metze and Shinji Watanabe

*SPLAT: Speech-Language Joint Pre-Training for Spoken Language Understanding*

Yu-An Chung, Chenguang Zhu and Michael Zeng

*Worldly Wise (WoW) - Cross-Lingual Knowledge Fusion for Fact-based Visual Spoken-Question Answering*

Kiran Ramnath, Leda Sari, Mark Hasegawa-Johnson and Chang Yoo

*Align-Refine: Non-Autoregressive Speech Recognition via Iterative Realignment*

Ethan A. Chi, Julian Salazar and Katrin Kirchhoff

**19:40–21:00 6B: NLP Applications**

*Everything Has a Cause: Leveraging Causal Inference in Legal Text Analysis*

Xiao Liu, Da Yin, Yansong Feng, Yuting Wu and Dongyan Zhao

*Counterfactual Supporting Facts Extraction for Explainable Medical Record Based Diagnosis with Graph Network*

Haoran Wu, Wei Chen, Shuang Xu and Bo Xu

*Personalized Response Generation via Generative Split Memory Network*

Yuwei Wu, Xuezhe Ma and Diyi Yang

*Towards Few-shot Fact-Checking via Perplexity*

Nayeon Lee, Yejin Bang, Andrea Madotto and Pascale Fung

**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

*Active<sup>2</sup> Learning: Actively reducing redundancies in Active Learning methods for Sequence Tagging and Machine Translation*

Rishi Hazra, Parag Dutta, Shubham Gupta, Mohammed Abdul Qaathir and Ambedkar Dukkipati

*Generating An Optimal Interview Question Plan Using A Knowledge Graph And Integer Linear Programming*

Soham Datta, Prabir Mallick, Sangameshwar Patil, Indrajit Bhattacharya and Girish Palshikar

**19:40–21:00 6C: Machine Learning for NLP: Classification and Structured Prediction Models**

*Model Extraction and Adversarial Transferability, Your BERT is Vulnerable!*

Xuanli He, Lingjuan Lyu, Lichao Sun and Qiongkai Xu

*A Global Past-Future Early Exit Method for Accelerating Inference of Pre-trained Language Models*

Kaiyuan Liao, Yi Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He

*Masked Conditional Random Fields for Sequence Labeling*

Tianwen Wei, Jianwei Qi, Shenghuan He and Songtao Sun

*Heterogeneous Graph Neural Networks for Concept Prerequisite Relation Learning in Educational Data*

Chenghao Jia, Yongliang Shen, Yechun Tang, Lu Sun and Weiming Lu

*Be Careful about Poisoned Word Embeddings: Exploring the Vulnerability of the Embedding Layers in NLP Models*

Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun and Bin He

*DA-Transformer: Distance-aware Transformer*

Chuhan Wu, Fangzhao Wu and Yongfeng Huang

Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)

**19:40–21:00 6D: Language Resources and Evaluation**

*ASAP: A Chinese Review Dataset Towards Aspect Category Sentiment Analysis and Rating Prediction*

Jiahao Bu, Lei Ren, Shuang Zheng, Yang Yang, Jingang Wang, Fuzheng Zhang and Wei Wu

*Are NLP Models really able to Solve Simple Math Word Problems?*

Arkil Patel, Satwik Bhattamishra and Navin Goyal

*WRIME: A New Dataset for Emotional Intensity Estimation with Subjective and Objective Annotations*

Tomoyuki Kajiwara, Chenhui Chu, Noriko Takemura, Yuta Nakashima and Hajime Nagahara

*KPQA: A Metric for Generative Question Answering Using Keyphrase Weights*

Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Joongbo Shin and Kyomin Jung

*StylePTB: A Compositional Benchmark for Fine-grained Controllable Text Style Transfer*

Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov and Louis-Philippe Morency

*Blow the Dog Whistle: A Chinese Dataset for Cant Understanding with Common Sense and World Knowledge*

Canwen Xu, Wangchunshu Zhou, Tao Ge, Ke Xu, Julian McAuley and Furu Wei

*COVID-19 Named Entity Recognition for Vietnamese*

Thinh Hung Truong, Mai Hoang Dao and Dat Quoc Nguyen



**Mon 07 Jun 2021 (all times PDT, UTC-7) (continued)**

**19:40–21:00 6E: Computational Social Science and Cultural Analytics**

*Framing Unpacked: A Semi-Supervised Interpretable Multi-View Model of Media Frames*

Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak, Andrew Turpin and Lea Frermann

*Automatic Classification of Neutralization Techniques in the Narrative of Climate Change Scepticism*

Shraey Bhatia, Jey Han Lau and Timothy Baldwin

*Suicide Ideation Detection via Social and Temporal User Representations using Hyperbolic Learning*

Ramit Sawhney, Harshit Joshi, Rajiv Ratn Shah and Lucie Flek

*WikiTalkEdit: A Dataset for modeling Editors' behaviors on Wikipedia*

Kokil Jaidka, Andrea Ceolin, Iknoor Singh, Niyati Chhaya and Lyle Ungar

*The structure of online social networks modulates the rate of lexical change*

Jian Zhu and David Jurgens

*Modeling Framing in Immigration Discourse on Social Media*

Julia Mendelsohn, Ceren Budak and David Jurgens

**Tue 08 Jun 2021 (all times PDT, UTC-7)**

**08:00–09:00 Keynote**

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**09:00–10:20 7A: Computational Social Science and Cultural Analytics**

*Modeling the Severity of Complaints in Social Media*

Mali Jin and Nikolaos Aletras

*What About the Precedent: An Information-Theoretic Analysis of Common Law*

Josef Valvoda, Tiago Pimentel, Niklas Stoehr, Ryan Cotterell and Simone Teufel

*Introducing CAD: the Contextual Abuse Dataset*

Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini and Rebekah Tromble

*Lifelong Learning of Hate Speech Classification on Social Media*

Jing Qian, Hong Wang, Mai ElSherief and Xifeng Yan

*Learning to Recognize Dialect Features*

Dorottya Demszky, Devyani Sharma, Jonathan Clark, Vinodkumar Prabhakaran and Jacob Eisenstein

*[TACL15] Characterizing English Variation across Social Media Communities with BERT*

Lucy Li, David Bamman

**09:00–10:20 7B: Green NLP**

*It's Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners*

Timo Schick and Hinrich Schütze

*Static Embeddings as Efficient Knowledge Bases?*

Philipp Dufter, Nora Kassner and Hinrich Schütze

*Highly Efficient Knowledge Graph Embedding Learning with Orthogonal Procrustes Analysis*

Xutan Peng, Guanyi Chen, Chenghua Lin and Mark Stevenson

*Rethinking Network Pruning – under the Pre-train and Fine-tune Paradigm*

Dongkuan Xu, Ian En-Hsu Yen, Jinxi Zhao and Zhibin Xiao

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Towards a Comprehensive Understanding and Accurate Evaluation of Societal Biases in Pre-Trained Transformers*

Andrew Silva, Pradyumna Tambwekar and Matthew Gombolay

*Detoxifying Language Models Risks Marginalizing Minority Voices*

Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap and Dan Klein

*HONEST: Measuring Hurtful Sentence Completion in Language Models*

Debora Nozza, Federico Bianchi and Dirk Hovy

**09:00–10:20 7C: Language Grounding to Vision, Robotics and Beyond**

*EaSe: A Diagnostic Tool for VQA based on Answer Diversity*

Shailza Jolly, Sandro Pezzelle and Moin Nabi

*DeCEMBERT: Learning from Noisy Instructional Videos via Dense Captions and Entropy Minimization*

Zineng Tang, Jie Lei and Mohit Bansal

*Improving Generation and Evaluation of Visual Stories via Semantic Consistency*

Adyasha Maharana, Darryl Hannan and Mohit Bansal

*Multilingual Multimodal Pre-training for Zero-Shot Cross-Lingual Transfer of Vision-Language Models*

Po-Yao Huang, Mandela Patrick, Junjie Hu, Graham Neubig, Florian Metze and Alexander Hauptmann

*Video Question Answering with Phrases via Semantic Roles*

Arka Sadhu, Kan Chen and Ram Nevatia

*[TACL10] Latent Compositional Representations Improve Systematic Generalization in Grounded Question Answering*

Ben Bogin: ben.bogin@, Jonathan Berant, Sanjay Subramanian, Matt Gardner

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**09:00–10:20 7D: Language Resources and Evaluation**

*From Masked Language Modeling to Translation: Non-English Auxiliary Tasks Improve Zero-shot Spoken Language Understanding*

Rob van der Goot, Ibrahim Sharaf, Aizhan Imankulova, Ahmet Üstün, Marija Stepanović, Alan Ramponi, Siti Oryza Khairunnisa, Mamoru Komachi and Barbara Plank

*WEC: Deriving a Large-scale Cross-document Event Coreference dataset from Wikipedia*

Alon Eirew, Arie Cattan and Ido Dagan

*Challenging distributional models with a conceptual network of philosophical terms*

Yvette Oortwijn, Jelke Bloem, Pia Sommerauer, Francois Meyer, Wei Zhou and Antske Fokkens

*KILT: a Benchmark for Knowledge Intensive Language Tasks*

Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel and Sebastian Riedel

*[TACL3] AMR Similarity Metrics from Principles*

Juri Opitz, Letitia Parcalabescu, Anette Frank

*[TACL19] Evaluating Document Coherence Modelling*

Aili Shen, Meladel Mistica, Bahar Salehi, Hang Li, Timothy Baldwin, Jianzhong Qi

**09:00–10:20 7E: Machine Learning for NLP: Classification and Structured Prediction Models**

*A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios*

Michael A. Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen and Dietrich Klakow

*Temporal Knowledge Graph Completion using a Linear Temporal Regularizer and Multivector Embeddings*

Chengjin Xu, Yung-Yu Chen, Mojtaba Nayyeri and Jens Lehmann

*UDALM: Unsupervised Domain Adaptation through Language Modeling*

Constantinos Karouzos, Georgios Paraskevopoulos and Alexandros Potamianos

*Beyond Black & White: Leveraging Annotator Disagreement via Soft-Label Multi-Task Learning*

Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy and Massimo Poesio

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Clustering-based Inference for Biomedical Entity Linking*

Rico Angell, Nicholas Monath, Sunil Mohan, Nishant Yadav and Andrew McCallum

*Variance-reduced First-order Meta-learning for Natural Language Processing Tasks*

Lingxiao Wang, Kevin Huang, Tengyu Ma, Quanquan Gu and Jing Huang

*Diversity-Aware Batch Active Learning for Dependency Parsing*

Tianze Shi, Adrian Benton, Igor Malioutov and Ozan Irsoy

**10:20–11:40 8A: Machine Learning for NLP: Language Modeling and Sequence to Sequence Models**

*How many data points is a prompt worth?*

Teven Le Scao and Alexander Rush

*Can Latent Alignments Improve Autoregressive Machine Translation?*

Adi Haviv, Lior Vassertail and Omer Levy

*Smoothing and Shrinking the Sparse Seq2Seq Search Space*

Ben Peters and André F. T. Martins

*Unified Pre-training for Program Understanding and Generation*

Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray and Kai-Wei Chang

*Hyperparameter-free Continuous Learning for Domain Classification in Natural Language Understanding*

Ting Hua, Yilin Shen, Changsheng Zhao, Yen-Chang Hsu and Hongxia Jin

*[TACL5] A Primer in BERTology: What We Know About How BERT Works*

Anna Rogers, Olga Kovaleva, Anna Rumshisky

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**10:20–11:40 8B: NLP Applications**

*On the Embeddings of Variables in Recurrent Neural Networks for Source Code*

Nadezhda Chirkova

*Cross-Lingual Word Embedding Refinement by  $\ell_1$  Norm Optimisation*

Xutan Peng, Chenghua Lin and Mark Stevenson

*Semantic Frame Forecast*

Chieh-Yang Huang and Ting-Hao Huang

*MUSER: MULTImodal Stress detection using Emotion Recognition as an Auxiliary Task*

Yiqun Yao, Michalis Papakostas, Mihai Burzo, Mohamed Abouelenien and Rada Mihalcea

*Learning to Decompose and Organize Complex Tasks*

Yi Zhang, Sujay Kumar Jauhar, Julia Kiseleva, Ryen White and Dan Roth

*Continual Learning for Text Classification with Information Disentanglement Based Regularization*

Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang and Diyi Yang

**10:20–11:40 8C: Sentence-level Semantics and Textual Inference**

*Learning from Executions for Semantic Parsing*

Bailin Wang, Mirella Lapata and Ivan Titov

*Learning to Synthesize Data for Semantic Parsing*

Bailin Wang, Wenpeng Yin, Xi Victoria Lin and Caiming Xiong

*Edge: Enriching Knowledge Graph Embeddings with External Text*

Saed Rezayi, Handong Zhao, Sungchul Kim, Ryan Rossi, Nedim Lipka and Sheng Li

*FLIN: A Flexible Natural Language Interface for Web Navigation*

Sahisnu Mazumder and Oriana Riva

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Game-theoretic Vocabulary Selection via the Shapley Value and Banzhaf Index*

Roma Patel, Marta Garnelo, Ian Gemp, Chris Dyer and Yoram Bachrach

*Incorporating External Knowledge to Enhance Tabular Reasoning*

J. Neeraja, Vivek Gupta and Vivek Srikumar

*Compositional Generalization for Neural Semantic Parsing via Span-level Supervised Attention*

Pengcheng Yin, Hao Fang, Graham Neubig, Adam Pauls, Emmanouil Antonios Platanios, Yu Su, Sam Thomson and Jacob Andreas

**10:20–11:40 8D: Sentiment Analysis and Stylistic Analysis**

*Domain Adaptation for Arabic Cross-Domain and Cross-Dialect Sentiment Analysis from Contextualized Word Embedding*

Abdellah El Mekki, Abdelkader El Mahdaoui, Ismail Berrada and Ahmed Khoumsi

*Multi-task Learning of Negation and Speculation for Targeted Sentiment Classification*

Andrew Moore and Jeremy Barnes

*A Disentangled Adversarial Neural Topic Model for Separating Opinions from Plots in User Reviews*

Gabriele Pergola, Lin Gui and Yulan He

*Graph Ensemble Learning over Multiple Dependency Trees for Aspect-level Sentiment Classification*

Xiaochen Hou, Peng Qi, Guangtao Wang, Rex Ying, Jing Huang, Xiaodong He and Bowen Zhou

*Emotion-Infused Models for Explainable Psychological Stress Detection*

Elsbeth Turcan, Smaranda Muresan and Kathleen McKeown

*Aspect-based Sentiment Analysis with Type-aware Graph Convolutional Networks and Layer Ensemble*

Yuanhe Tian, Guimin Chen and Yan Song

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**10:20–11:40 8E: Syntax: Tagging, Chunking, and Parsing**

*Supertagging-based Parsing with Linear Context-free Rewriting Systems*

Thomas Ruprecht and Richard Mörbitz

*Outside Computation with Superior Functions*

Parker Riley and Daniel Gildea

*Learning Syntax from Naturally-Occurring Bracketings*

Tianze Shi, Ozan İrsoy, Igor Malioutov and Lillian Lee

*[CL1] What Should/Do/Can LSTMs Learn When Parsing Auxiliary Verb Constructions?*

Miryam de Lhoneux, Sara Stymne, Joakim Nivre

*[TACL2] Reducing Confusion in Active Learning for Part-Of-Speech Tagging*

Aditi Chaudhary, Antonios Anastasopoulos, Zaid Sheikh, Graham Neubig

**10:20–11:40 Business Meeting**

**17:00–18:20 9A: Dialogue and Interactive Systems**

*Bot-Adversarial Dialogue for Safe Conversational Agents*

Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston and Emily Dinan

*Non-Autoregressive Semantic Parsing for Compositional Task-Oriented Dialog*

Arun Babu, Akshat Shrivastava, Armen Aghajanyan, Ahmed Aly, Angela Fan and Marjan Ghazvininejad

*Example-Driven Intent Prediction with Observers*

Shikib Mehri and Mihail Eric

*Imperfect also Deserves Reward: Multi-Level and Sequential Reward Modeling for Better Dialog Management*

Zhengxu Hou, Bang Liu, Ruihui Zhao, Zijing Ou, Yafei Liu, Xi Chen and Yefeng Zheng

*Action-Based Conversations Dataset: A Corpus for Building More In-Depth Task-Oriented Dialogue Systems*

Derek Chen, Howard Chen, Yi Yang, Alexander Lin and Zhou Yu



**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Controlling Dialogue Generation with Semantic Exemplars*  
Prakhar Gupta, Jeffrey Bigham, Yulia Tsvetkov and Amy Pavel

**17:00–18:20 9B: Information Retrieval and Text Mining**

*COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List*  
Luyu Gao, Zhuyun Dai and Jamie Callan

*X-Class: Text Classification with Extremely Weak Supervision*  
Zihan Wang, Dheeraj Mekala and Jingbo Shang

*Fine-tuning Encoders for Improved Monolingual and Zero-shot Polylingual Neural Topic Modeling*  
Aaron Mueller and Mark Dredze

*Exploring the Relationship Between Algorithm Performance, Vocabulary, and Run-Time in Text Classification*  
Wilson Fearn, Orion Weller and Kevin Seppi

*Faithfully Explainable Recommendation via Neural Logic Reasoning*  
Yaxin Zhu, Yikun Xian, Zuohui Fu, Gerard de Melo and Yongfeng Zhang

*You Sound Like Someone Who Watches Drama Movies: Towards Predicting Movie Preferences from Conversational Interactions*  
Sergey Volokhin, Joyce Ho, Oleg Rokhlenko and Eugene Agichtein

*[TACL8] Sparse, Dense, and Attentional Representations for Text Retrieval*  
Yi Luan, Jacob Eisenstein, Kristina Toutanova, Michael Collins

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**17:00–18:20 9C: Language Grounding to Vision, Robotics and Beyond**

*Reading and Acting while Blindfolded: The Need for Semantics in Text Game Agents*

Shunyu Yao, Karthik Narasimhan and Matthew Hausknecht

*SOrT-ing VQA Models : Contrastive Gradient Learning for Improved Consistency*

Sameer Dharur, Purva Tendulkar, Dhruv Batra, Devi Parikh and Ramprasaath R. Selvaraju

*Semi-Supervised Policy Initialization for Playing Games with Language Hints*

Tsu-Jui Fu and William Yang Wang

*Revisiting Document Representations for Large-Scale Zero-Shot Learning*

Jihyung Kil and Wei-Lun Chao

**17:00–18:20 9D: Language Resources and Evaluation**

*Negative language transfer in learner English: A new dataset*

Leticia Farias Wanderley, Nicole Zhao and Carrie Demmans Epp

*SentSim: Crosslingual Semantic Evaluation of Machine Translation*

Yurun Song, Junchen Zhao and Lucia Specia

*Quality Estimation for Image Captions Based on Large-scale Human Evaluations*

Tomer Levinboim, Ashish V. Thapliyal, Piyush Sharma and Radu Soricut

*CaSiNo: A Corpus of Campsite Negotiation Dialogues for Automatic Negotiation Systems*

Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May and Jonathan Gratch

*News Headline Grouping as a Challenging NLU Task*

Philippe Laban, Lucas Bandarkar and Marti A. Hearst

*Olá, Bonjour, Salve! XFORMAL: A Benchmark for Multilingual Formality Style Transfer*

Eleftheria Briakou, Di Lu, Ke Zhang and Joel Tetreault

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**17:00–18:20 9E: Machine Learning for NLP: Classification and Structured Prediction Models**

*Grouping Words with Semantic Diversity*

Karine Chubarian, Abdul Rafae Khan, Anastasios Sidiropoulos and Jia Xu

*Noise Stability Regularization for Improving BERT Fine-tuning*

Hang Hua, Xingjian Li, Dejing Dou, Chengzhong Xu and Jiebo Luo

*FlowPrior: Learning Expressive Priors for Latent Variable Sentence Models*

Xiaoan Ding and Kevin Gimpel

*HTCInfoMax: A Global Model for Hierarchical Text Classification via Information Maximization*

Zhongfen Deng, Hao Peng, Dongxiao He, Jianxin Li and Philip Yu

*[TACL7] Modeling Content and Context with Deep Relational Learning*

Maria Leonor Pacheco, Dan Goldwasser

*Knowledge Guided Metric Learning for Few-Shot Text Classification*

Dianbo Sui, Yubo Chen, Binjie Mao, Delai Qiu, Kang Liu and Jun Zhao

**18:20–19:40 10A: Dialogue and Interactive Systems**

*Ensemble of MRR and NDCG models for Visual Dialog*

Idan Schwartz

*Supervised Neural Clustering via Latent Structured Output Learning: Application to Question Intents*

Iryna Haponchyk and Alessandro Moschitti

*ConVEx: Data-Efficient and Few-Shot Slot Labeling*

Matthew Henderson and Ivan Vulić

*CREAD: Combined Resolution of Ellipses and Anaphora in Dialogues*

Bo-Hsiang Tseng, Shruti Bhargava, Jiarui Lu, Joel Ruben Antony Moniz, Dhivya Piraviperumal, Lin Li and Hong Yu

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Knowledge-Driven Slot Constraints for Goal-Oriented Dialogue Systems*

Piyawat Lertvittayakumjorn, Daniele Bonadiman and Saab Mansour

*Clipping Loops for Sample-Efficient Dialogue Policy Optimisation*

Yen-Chen Wu and Carl Edward Rasmussen

**18:20–19:40 10B: Information Extraction**

*Integrating Lexical Information into Entity Neighbourhood Representations for Relation Prediction*

Ian Wood, Mark Johnson and Stephen Wan

*Noisy-Labeled NER with Confidence Estimation*

Kun Liu, Yao Fu, Chuanqi Tan, Mosha Chen, Ningyu Zhang, Songfang Huang and Sheng Gao

*TABBIE: Pretrained Representations of Tabular Data*

Hiroshi Iida, Dung Thai, Varun Manjunatha and Mohit Iyyer

*Better Feature Integration for Named Entity Recognition*

Lu Xu, Zhanming Jie, Wei Lu and Lidong Bing

*ZS-BERT: Towards Zero-Shot Relation Extraction with Attribute Representation Learning*

Chih-Yao Chen and Cheng-Te Li

*Graph Convolutional Networks for Event Causality Identification with Rich Document-level Structures*

Minh Tran Phu and Thien Huu Nguyen

*A Context-Dependent Gated Module for Incorporating Symbolic Semantics into Event Coreference Resolution*

Tuan Lai, Heng Ji, Trung Bui, Quan Hung Tran, Franck Dernoncourt and Walter Chang

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**18:20–19:40 10C: Language Generation**

*Multi-Style Transfer with Discriminative Feedback on Disjoint Corpus*

Navita Goyal, Balaji Vasan Srinivasan, Anandhavelu N and Abhilasha Sancheti

*FUDGE: Controlled Text Generation With Future Discriminators*

Kevin Yang and Dan Klein

*Controllable Text Simplification with Explicit Paraphrasing*

Mounica Maddela, Fernando Alva-Manchego and Wei Xu

*Knowledge Graph Based Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-training*

Oshin Agarwal, Heming Ge, Siamak Shakeri and Rami Al-Rfou

*Choose Your Own Adventure: Paired Suggestions in Collaborative Writing for Evaluating Story Generation Models*

Elizabeth Clark and Noah A. Smith

*[TACL17] There Once Was a Really Bad Poet, It Was Automated but You Didn't Know It*

Jianyong, jw542@duke.edu, Xiaoxuan, zhangxiaoxuanaa@gmail.com, Yuren Zhou, Christopher Suh, Cynthia Rudin

**18:20–19:40 10D: Multilinguality**

*InfoXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training*

Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang and Ming Zhou

*Context-Interactive Pre-Training for Document Machine Translation*

Pengcheng Yang, Pei Zhang, Boxing Chen, Jun Xie and Weihua Luo

*Code-Mixing on Sesame Street: Dawn of the Adversarial Polyglots*

Samson Tan and Shafiq Joty

*X-METRA-ADA: Cross-lingual Meta-Transfer learning Adaptation to Natural Language Understanding and Question Answering*

Meryem M'hamdi, Doo Soon Kim, Franck Dernoncourt, Trung Bui, Xiang Ren and Jonathan May

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Explicit Alignment Objectives for Multilingual Bidirectional Encoders*

Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant and Graham Neubig

*Cross-lingual Cross-modal Pretraining for Multimodal Retrieval*

Hongliang Fei, Tan Yu and Ping Li

*Wikipedia Entities as Rendezvous across Languages: Grounding Multilingual Language Models by Predicting Wikipedia Hyperlinks*

Iacer Calixto, Alessandro Raganato and Tommaso Pasini

**18:20–19:40 10E: Question Answering**

*multiPProver: Generating Multiple Proofs for Improved Interpretability in Rule Reasoning*

Swarnadeep Saha, Prateek Yadav and Mohit Bansal

*Adaptable and Interpretable Neural Memory Over Symbolic Knowledge*

Pat Verga, Haitian Sun, Livio Baldini Soares and William Cohen

*CLEVR\_HYP: A Challenge Dataset and Baselines for Visual Question Answering with Hypothetical Actions over Images*

Shailaja Keyur Sampat, Akshay Kumar, Yezhou Yang and Chitta Baral

*Refining Targeted Syntactic Evaluation of Language Models*

Benjamin Newman, Kai-Siang Ang, Julia Gong and John Hewitt

*Universal Adversarial Attacks with Natural Triggers for Text Classification*

Liwei Song, Xinwei Yu, Hsuan-Tung Peng and Karthik Narasimhan

*QuadrupletBERT: An Efficient Model For Embedding-Based Large-Scale Retrieval*

Peiyang Liu, Sen Wang, Xi Wang, Wei Ye and Shikun Zhang

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**19:40–21:00 11A: Ethics, Bias, and Fairness**

*Dynamically Disentangling Social Bias from Task-Oriented Representations with Adversarial Attack*

Liwen Wang, Yuanmeng Yan, Keqing He, Yanan Wu and Weiran Xu

*An Empirical Investigation of Bias in the Multimodal Analysis of Financial Earnings Calls*

Ramit Sawhney, Arshiya Aggarwal and Rajiv Ratn Shah

*Beyond Fair Pay: Ethical Implications of NLP Crowdsourcing*

Boaz Shmueli, Jan Fell, Soumya Ray and Lun-Wei Ku

*On Transferability of Bias Mitigation Effects in Language Model Fine-Tuning*

Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves and Xiang Ren

*Case Study: Deontological Ethics in NLP*

Shrimai Prabhumoye, Brendon Boldt, Ruslan Salakhutdinov and Alan W Black

*Privacy Regularization: Joint Privacy-Utility Optimization in Language Models*

Fatemehsadat Mireshghallah, Huseyin Inan, Marcello Hasegawa, Victor Rühle, Taylor Berg-Kirkpatrick and Robert Sim

*On the Impact of Random Seeds on the Fairness of Clinical Classifiers*

Silvio Amir, Jan-Willem van de Meent and Byron Wallace

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**19:40–21:00 11B: Interpretability and Analysis of Models for NLP**

*Topic Model or Topic Twaddle? Re-evaluating Semantic Interpretability Measures*

Caitlin Doogan and Wray Buntine

*Discourse Probing of Pretrained Language Models*

Fajri Koto, Jey Han Lau and Timothy Baldwin

*UniDrop: A Simple yet Effective Technique to Improve Transformer without Extra Cost*

Zhen Wu, Lijun Wu, Qi Meng, Yingce Xia, Shufang Xie, Tao Qin, Xinyu Dai and Tie-Yan Liu

*tWT-WT: A Dataset to Assert the Role of Target Entities for Detecting Stance of Tweets*

Ayush Kaushal, Avirup Saha and Niloy Ganguly

*Learning to Learn to be Right for the Right Reasons*

Pride Kavumba, Benjamin Heinzerling, Ana Brassard and Kentaro Inui

*Double Perturbation: On the Robustness of Robustness and Counterfactual Bias Evaluation*

Chong Zhang, Jieyu Zhao, Huan Zhang, Kai-Wei Chang and Cho-Jui Hsieh

*Explaining Neural Network Predictions on Sentence Pairs via Learning Word-Group Masks*

Hanjie Chen, Song Feng, Jatin Ganhotra, Hui Wan, Chulaka Gunasekara, Sachindra Joshi and Yangfeng Ji



**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**19:40–21:00 11C: Machine Translation**

*Almost Free Semantic Draft for Neural Machine Translation*

Xi Ai and Bin Fang

*Pruning-then-Expanding Model for Domain Adaptation of Neural Machine Translation*

Shuhao Gu, Yang Feng and Wanying Xie

*Multi-Hop Transformer for Document-Level Machine Translation*

Long Zhang, Tong Zhang, Haibo Zhang, Baosong Yang, Wei Ye and Shikun Zhang

*Continual Learning for Neural Machine Translation*

Yue Cao, Hao-Ran Wei, Boxing Chen and Xiaojun Wan

*Self-Training for Unsupervised Neural Machine Translation in Unbalanced Training Data Scenarios*

Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita and Tiejun Zhao

*Smart-Start Decoding for Neural Machine Translation*

Jian Yang, Shuming Ma, Dongdong Zhang, Juncheng Wan, Zhoujun Li and Ming Zhou

*Multi-Task Learning with Shared Encoder for Non-Autoregressive Machine Translation*

Yongchang Hao, Shilin He, Wenxiang Jiao, Zhaopeng Tu, Michael Lyu and Xing Wang

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

**19:40–21:00 11D: NLP Applications**

*ER-AE: Differentially Private Text Generation for Authorship Anonymization*

Haohan Bo, Steven H. H. Ding, Benjamin C. M. Fung and Farkhund Iqbal

*Distantly Supervised Transformers For E-Commerce Product QA*

Happy Mittal, Aniket Chakrabarti, Belhassen Bayar, Animesh Anant Sharma and Nikhil Rasiwasia

*Quantitative Day Trading from Natural Language using Reinforcement Learning*

Ramit Sawhney, Arnav Wadhwa, Shivam Agarwal and Rajiv Ratn Shah

*Restoring and Mining the Records of the Joseon Dynasty via Neural Language Modeling and Machine Translation*

Kyeongpil Kang, Kyohoon Jin, Soyoung Yang, Soojin Jang, Jaegul Choo and Youngbin Kim

*Modeling Diagnostic Label Correlation for Automatic ICD Coding*

Shang-Chi Tsai, Chao-Wei Huang and Yun-Nung Chen

*Self-Supervised Contrastive Learning for Efficient User Satisfaction Prediction in Conversational Agents*

Mohammad Kachuee, Hao Yuan, Young-Bum Kim and Sungjin Lee

**19:40–21:00 11E: Special Theme: New Challenges in NLP**

*A recipe for annotating grounded clarifications*

Luciana Benotti and Patrick Blackburn

*Grey-box Adversarial Attack And Defence For Sentiment Classification*

Ying Xu, Xu Zhong, Antonio Jimeno Yepes and Jey Han Lau

*How low is too low? A monolingual take on lemmatisation in Indian languages*

Kumar Saunack, Kumar Saurav and Pushpak Bhattacharyya

*Causal Effects of Linguistic Properties*

Reid Pryzant, Dallas Card, Dan Jurafsky, Victor Veitch and Dhanya Sridhar

**Tue 08 Jun 2021 (all times PDT, UTC-7) (continued)**

*Dynabench: Rethinking Benchmarking in NLP*

Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts and Adina Williams

*Translational NLP: A New Paradigm and General Principles for Natural Language Processing Research*

Denis Newman-Griffis, Jill Fain Lehman, Carolyn Rosé and Harry Hochheiser

**Wed 09 Jun 2021 (all times PDT, UTC-7)**

**08:00–09:00** Keynote

**09:00–10:20** 12A: Discourse and Pragmatics

*Predicting Discourse Trees from Transformer-based Neural Summarizers*

Wen Xiao, Patrick Huber and Giuseppe Carenini

*Probing for Bridging Inference in Transformer Language Models*

Onkar Pandit and Yufang Hou

*Is Incoherence Surprising? Targeted Evaluation of Coherence Prediction from Language Models*

Anne Beyer, Sharid Loáiciga and David Schlangen

*Stay Together: A System for Single and Split-antecedent Anaphora Resolution*

Juntao Yu, Nafise Sadat Moosavi, Silviu Paun and Massimo Poesio

*[TACL13] Decontextualization: Making Sentences Stand-Alone*

Eunsol Choi, Jennimaria Palomaki, Matthew Lamm, Tom Kwiatkowski, Dipanjan Das, Michael Collins

*[CL4] Universal Discourse Representation Structure Parsing*

Jiangming Liu, Shay B. Cohen, Mirella Lapata, Johan Bos

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**09:00–10:20 12B: Information Retrieval and Text Mining**

*Redefining Absent Keyphrases and their Effect on Retrieval Effectiveness*

Florian Boudin and Ygor Gallina

*CoRT: Complementary Rankings from Transformers*

Marco Wrzalik and Dirk Krechel

*Multi-source Neural Topic Modeling in Multi-view Embedding Spaces*

Pankaj Gupta, Yatin Chaudhary and Hinrich Schütze

*Inductive Topic Variational Graph Auto-Encoder for Text Classification*

Qianqian Xie, Jimin Huang, Pan Du, Min Peng and Jian-Yun Nie

*Self-Alignment Pretraining for Biomedical Entity Representations*

Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella and Nigel Collier

*TaxoClass: Hierarchical Multi-Label Text Classification Using Only Class Names*

Jiaming Shen, Wenda Qiu, Yu Meng, Jingbo Shang, Xiang Ren and Jiawei Han

**09:00–10:20 12C: Language Generation**

*MERMAID: Metaphor Generation with Symbolism and Discriminative Decoding*

Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan and Nanyun Peng

*On Learning Text Style Transfer with Direct Rewards*

Yixin Liu, Graham Neubig and John Wieting

*Focused Attention Improves Document-Grounded Generation*

Shrimai Prabhumoye, Kazuma Hashimoto, Yingbo Zhou, Alan W Black and Ruslan Salakhutdinov

*NeuroLogic Decoding: (Un)supervised Neural Text Generation with Predicate Logic Constraints*

Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula and Yejin Choi

**Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)**

*Ask what's missing and what's useful: Improving Clarification Question Generation using Global Knowledge*

Bodhisattwa Prasad Majumder, Sudha Rao, Michel Galley and Julian McAuley

*Progressive Generation of Long Text with Pretrained Language Models*

Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric Xing and Zhiting Hu

**09:00–10:20 12D: Language Resources and Evaluation**

*SOCER: An Information-Sparse Discourse State Tracking Collection in the Sports Commentary Domain*

Ruo Chen Zhang and Carsten Eickhoff

*Plot-guided Adversarial Example Construction for Evaluating Open-domain Story Generation*

Sarik Ghazarian, Zixi Liu, Akash S M, Ralph Weischedel, Aram Galstyan and Nanyun Peng

*MultiOpEd: A Corpus of Multi-Perspective News Editorials*

Siyi Liu, Sihao Chen, Xander Uyttendaele and Dan Roth

*Swords: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality*

Mina Lee, Chris Donahue, Robin Jia, Alexander Iyabor and Percy Liang

*"I'm Not Mad": Commonsense Implications of Negation and Contradiction*

Liwei Jiang, Antoine Bosselut, Chandra Bhagavatula and Yejin Choi

*Identifying Medical Self-Disclosure in Online Communities*

Mina Valizadeh, Pardis Ranjbar-Noiey, Cornelia Caragea and Natalie Parde

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

09:00–10:20 12E: Linguistic Theories, Cognitive Modeling and Psycholinguistics

*Language in a (Search) Box: Grounding Language Learning in Real-World Human-Machine Interaction*

Federico Bianchi, Ciro Greco and Jacopo Tagliabue

*Finding Concept-specific Biases in Form–Meaning Associations*

Tiago Pimentel, Brian Roark, Søren Wichmann, Ryan Cotterell and Damián Blasi

*How (Non-)Optimal is the Lexicon?*

Tiago Pimentel, Irene Nikkarinen, Kyle Mahowald, Ryan Cotterell and Damián Blasi

*Word Complexity is in the Eye of the Beholder*

Sian Gooding, Ekaterina Kochmar, Seid Muhie Yimam and Chris Biemann

*Linguistic Complexity Loss in Text-Based Therapy*

Jason Wei, Kelly Finn, Emma Templeton, Thalia Wheatley and Soroush Vosoughi

*Ab Antiquo: Neural Proto-language Reconstruction*

Carlo Meloni, Shauli Ravfogel and Yoav Goldberg

*On Biasing Transformer Attention Towards Monotonicity*

Annette Rios, Chantal Amrhein, Noëmi Aepli and Rico Sennrich

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

10:20–11:40 13A: NLP Applications

*Extracting a Knowledge Base of Mechanisms from COVID-19 Papers*

Tom Hope, Aida Amini, David Wadden, Madeleine van Zuylen, Sravanthi Parasa, Eric Horvitz, Daniel Weld, Roy Schwartz and Hannaneh Hajishirzi

*Constrained Multi-Task Learning for Event Coreference Resolution*

Jing Lu and Vincent Ng

*Empirical Evaluation of Pre-trained Transformers for Human-Level NLP: The Role of Sample Size and Dimensionality*

Adithya V Ganesan, Matthew Matero, Aravind Reddy Ravula, Huy Vu and H. Andrew Schwartz

*Leveraging Deep Representations of Radiology Reports in Survival Analysis for Predicting Heart Failure Patient Mortality*

Hyun Gi Lee, Evan Sholle, Ashley Beecy, Subhi Al'Aref and Yifan Peng

*On the Use of Context for Predicting Citation Worthiness of Sentences in Scholarly Articles*

Rakesh Gosangi, Ravneet Arora, Mohsen Gheisarieha, Debanjan Mahata and Haimin Zhang

*Data and Model Distillation as a Solution for Domain-transferable Fact Verification*

Mitch Paul Mithun, Sandeep Suntwal and Mihai Surdeanu

*Adapting Coreference Resolution for Processing Violent Death Narratives*

Ankith Uppunda, Susan Cochran, Jacob Foster, Alina Arseniev-Koehler, Vickie Mays and Kai-Wei Chang

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

10:20–11:40 13B: Question Answering

*Time-Stamped Language Model: Teaching Language Models to Understand The Flow of Events*

Hossein Rajaby Faghihi and Parisa Kordjamshidi

*If You Want to Go Far Go Together: Unsupervised Joint Candidate Evidence Retrieval for Multi-hop Question Answering*

Vikas Yadav, Steven Bethard and Mihai Surdeanu

*SPARTQA: A Textual Question Answering Benchmark for Spatial Reasoning*

Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang Ning and Parisa Kordjamshidi

*A Dataset of Information-Seeking Questions and Answers Anchored in Research Papers*

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith and Matt Gardner

*Differentiable Open-Ended Commonsense Reasoning*

Bill Yuchen Lin, Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Xiang Ren and William Cohen

*Does Structure Matter? Encoding Documents for Machine Reading Comprehension*

Hui Wan, Song Feng, Chulaka Gunasekara, Siva Sankalp Patel, Sachindra Joshi and Luis Lastras

*Multi-Step Reasoning Over Unstructured Text with Beam Dense Retrieval*

Chen Zhao, Chenyan Xiong, Jordan Boyd-Graber and Hal Daumé III



Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**10:20–11:40 13C: Lexical Semantics**

*Scalable and Interpretable Semantic Change Detection*

Syrielle Montariol, Matej Martinc and Lidia Pivovarova

*Scalar Adjective Identification and Multilingual Ranking*

Aina Garí Soler and Marianna Apidianaki

*ESC: Redesigning WSD with Extractive Sense Comprehension*

Edoardo Barba, Tommaso Pasini and Roberto Navigli

*Recent advances in neural metaphor processing: A linguistic, cognitive and social perspective*

Xiaoyu Tong, Ekaterina Shutova and Martha Lewis

*Constructing Taxonomies from Pretrained Language Models*

Catherine Chen, Kevin Lin and Dan Klein

*Event Representation with Sequential, Semi-Supervised Discrete Variables*

Mehdi Rezaee and Francis Ferraro

**10:20–11:40 13D: Sentiment Analysis and Stylistic Analysis**

*Seq2Emo: A Sequence to Multi-Label Emotion Classification Model*

Chenyang Huang, Amine Trabelsi, Xuebin Qin, Nawshad Farruque, Lili Mou and Osmar Zaiane

*Knowledge Enhanced Masked Language Model for Stance Detection*

Kornraphop Kawintiranon and Lisa Singh

*Learning Paralinguistic Features from Audiobooks through Style Voice Conversion*

Zakaria Aldeneh, Matthew Perez and Emily Mower Provost

*Adapting BERT for Continual Learning of a Sequence of Aspect Sentiment Classification Tasks*

Zixuan Ke, Hu Xu and Bing Liu

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

*Adversarial Learning for Zero-Shot Stance Detection on Social Media*

Emily Allaway, Malavika Srikanth and Kathleen McKeown

10:20–11:40 13E: Summarization

*Efficiently Summarizing Text and Graph Encodings of Multi-Document Clusters*

Ramakanth Pasunuru, Mengwen Liu, Mohit Bansal, Sujith Ravi and Markus Dreyer

*Enriching Transformers with Structured Tensor-Product Representations for Abstractive Summarization*

Yichen Jiang, Asli Celikyilmaz, Paul Smolensky, Paul Soulos, Sudha Rao, Hamid Palangi, Roland Fernandez, Caitlin Smith, Mohit Bansal and Jianfeng Gao

*What's in a Summary? Laying the Groundwork for Advances in Hospital-Course Summarization*

Griffin Adams, Emily Alsentzer, Mert Ketenci, Jason Zucker and Noémie Elhadad

*Understanding Factuality in Abstractive Summarization with FRANK: A Benchmark for Factuality Metrics*

Artidoro Pagnoni, Vidhisha Balachandran and Yulia Tsvetkov

*GSum: A General Framework for Guided Neural Abstractive Summarization*

Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang and Graham Neubig

*[TACL11] WikiAsp: A Dataset for Multi-domain Aspect-based Summarization*

Hiroaki Hayashi, Prashant Budania, Peng Wang, Chris Ackerson, Raj Neervannan, Graham Neubig

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

10:20–11:40 Best Paper Presentations

17:00–18:20 14A: Computational Social Science and Cultural Analytics

*Multitask Learning for Emotionally Analyzing Sexual Abuse Disclosures*

Ramit Sawhney, Puneet Mathur, Taru Jain, Akash Kumar Gautam and Rajiv Ratn Shah

*Self Promotion in US Congressional Tweets*

Jun Wang, Kelly Cui and Bei Yu

*Profiling of Intertextuality in Latin Literature Using Word Embeddings*

Patrick J. Burns, James Brofos, Kyle Li, Pramit Chaudhuri and Joseph P. Dexter

*Identifying inherent disagreement in natural language inference*

Xinliang Frederick Zhang and Marie-Catherine de Marneffe

*Modeling Human Mental States with an Entity-based Narrative Graph*

I-Ta Lee, Maria Leonor Pacheco and Dan Goldwasser

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**17:00–18:20 14B: Generation and Summarization**

*A Simple and Efficient Multi-Task Learning Approach for Conditioned Dialogue Generation*

Yan Zeng and Jian-Yun Nie

*Hurdles to Progress in Long-form Question Answering*

Kalpesh Krishna, Aurko Roy and Mohit Iyyer

*ENTRUST: Argument Reframing with Language Models and Entailment*

Tuhin Chakrabarty, Christopher Hidey and Smaranda Muresan

*Paragraph-level Simplification of Medical Texts*

Ashwin Devaraj, Iain Marshall, Byron Wallace and Junyi Jessy Li

*An Empirical Study on Neural Keyphrase Generation*

Rui Meng, Xingdi Yuan, Tong Wang, Sanqiang Zhao, Adam Trischler and Daqing He

*Attention Head Masking for Inference Time Content Selection in Abstractive Summarization*

Shuyang Cao and Lu Wang

**17:00–18:20 14C: Interpretability and Analysis of Models for NLP**

*Factual Probing Is [MASK]: Learning vs. Learning to Recall*

Zexuan Zhong, Dan Friedman and Danqi Chen

*Evaluating Saliency Methods for Neural Language Models*

Shuoyang Ding and Philipp Koehn

*Contextualized Perturbation for Textual Adversarial Attack*

Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun and Bill Dolan

*DirectProbe: Studying Representations without Classifiers*

Yichu Zhou and Vivek Srikumar

**Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)**

*Evaluating the Values of Sources in Transfer Learning*

Md Rizwan Parvez and Kai-Wei Chang

*Too Much in Common: Shifting of Embeddings in Transformer Language Models and its Implications*

Daniel Biś, Maksim Podkorytov and Xiuwen Liu

**17:00–18:20 14D: Machine Learning for NLP: Language Modeling and Sequence to Sequence Models**

*On the Inductive Bias of Masked Language Modeling: From Statistical to Syntactic Dependencies*

Tianyi Zhang and Tatsunori Hashimoto

*Limitations of Autoregressive Models and Their Alternatives*

Chu-Cheng Lin, Aaron Jaech, Xin Li, Matthew R. Gormley and Jason Eisner

*On the Transformer Growth for Progressive BERT Training*

Xiaotao Gu, Liyuan Liu, Hongkun Yu, Jing Li, Chen Chen and Jiawei Han

*Revisiting Simple Neural Probabilistic Language Models*

Simeng Sun and Mohit Iyyer

*ReadTwice: Reading Very Large Documents with Memories*

Yury Zemlyanskiy, Joshua Ainslie, Michiel de Jong, Philip Pham, Ilya Eckstein and Fei Sha

*SCRIPT: Self-Critic PreTraining of Transformers*

Erik Nijkamp, Bo Pang, Ying Nian Wu and Caiming Xiong

*Learning How to Ask: Querying LMs with Mixtures of Soft Prompts*

Guanghui Qin and Jason Eisner

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**17:00–18:20 14E: NLP Applications**

*Nutri-bullets Hybrid: Consensual Multi-document Summarization*

Darsh Shah, Lili Yu, Tao Lei and Regina Barzilay

*AVA: an Automatic eValuation Approach for Question Answering Systems*

Thuy Vu and Alessandro Moschitti

*SpanPredict: Extraction of Predictive Document Spans with Neural Attention*

Vivek Subramanian, Matthew Engelhard, Sam Berchuck, Lique Chen, Ricardo Henao and Lawrence Carin

*Text Editing by Command*

Felix Faltings, Michel Galley, Gerold Hintz, Chris Brockett, Chris Quirk, Jianfeng Gao and Bill Dolan

*A Deep Metric Learning Approach to Account Linking*

Aleem Khan, Elizabeth Fleming, Noah Schofield, Marcus Bishop and Nicholas Andrews

*Improving Factual Completeness and Consistency of Image-to-Text Radiology Report Generation*

Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz and Dan Jurafsky

**18:20–19:40 15A: Language Grounding to Vision, Robotics and Beyond**

*Multimodal End-to-End Sparse Model for Emotion Recognition*

Wenliang Dai, Samuel Cahyawijaya, Zihan Liu and Pascale Fung

*MIMOQA: Multimodal Input Multimodal Output Question Answering*

Hrituraj Singh, Anshul Nasery, Denil Mehta, Aishwarya Agarwal, Jatin Lamba and Balaji Vasanth Srinivasan

*OCID-Ref: A 3D Robotic Dataset With Embodied Language For Clutter Scene Grounding*

Ke-Jyun Wang, Yun-Hsuan Liu, Hung-Ting Su, Jen-Wei Wang, Yu-Siang Wang, Winston Hsu and Wen-Chin Chen

*Unsupervised Vision-and-Language Pre-training Without Parallel Images and Captions*

Liunian Harold Li, Haoxuan You, Zhecan Wang, Alireza Zareian, Shih-Fu Chang and Kai-Wei Chang

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

*Multitasking Inhibits Semantic Drift*

Athul Paul Jacob, Mike Lewis and Jacob Andreas

*Probing Contextual Language Models for Common Ground with Visual Representations*

Gabriel Ilharco, Rowan Zellers, Ali Farhadi and Hannaneh Hajishirzi

**18:20–19:40 15B: Machine Learning for NLP: Classification and Structured Prediction Models**

*BBAEG: Towards BERT-based Biomedical Adversarial Example Generation for Text Classification*

Ishani Mondal

*Targeted Adversarial Training for Natural Language Understanding*

Lis Pereira, Xiaodong Liu, Hao Cheng, Hoifung Poon, Jianfeng Gao and Ichiro Kobayashi

*Latent-Optimized Adversarial Neural Transfer for Sarcasm Detection*

Xu Guo, Boyang Li, Han Yu and Chunyan Miao

*Self-training Improves Pre-training for Natural Language Understanding*

Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Veselin Stoyanov and Alexis Conneau

*Supporting Clustering with Contrastive Learning*

Dejiao Zhang, Feng Nan, Xiaokai Wei, Shang-Wen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew O. Arnold and Bing Xiang

*[TACL16] Self-supervised Regularization for Text Classification*

Meng Zhou, Zechen Li, Pengtao Xie

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

18:20–19:40 15C: NLP Applications

*TITA: A Two-stage Interaction and Topic-Aware Text Matching Model*

Xingwu Sun, Yanling Cui, Hongyin Tang, Qiuyu Zhu, Fuzheng Zhang and Beihong Jin

*Neural Quality Estimation with Multiple Hypotheses for Grammatical Error Correction*

Zhenghao Liu, Xiaoyuan Yi, Maosong Sun, Liner Yang and Tat-Seng Chua

*Neural Network Surgery: Injecting Data Patterns into Pre-trained Models with Minimal Instance-wise Side Effects*

Zhiyuan Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He

*Discrete Argument Representation Learning for Interactive Argument Pair Identification*

Lu Ji, Zhongyu Wei, Jing Li, Qi Zhang and Xuanjing Huang

*On Unifying Misinformation Detection*

Nayeon Lee, Belinda Z. Li, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih and Madian Khabsa

*Frustratingly Easy Edit-based Linguistic Steganography with a Masked Language Model*

Honai Ueoka, Yugo Murawaki and Sadao Kurohashi

*Few-Shot Text Classification with Triplet Networks, Data Augmentation, and Curriculum Learning*

Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng and Shiqi Xu



Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**18:20–19:40 15D: Phonology, Morphology and Word Segmentation**

*Do RNN States Encode Abstract Phonological Alternations?*

Miikka Silfverberg, Francis Tyers, Garrett Nicolai and Mans Hulden

*Pre-training with Meta Learning for Chinese Word Segmentation*

Zhen Ke, Liang Shi, Songtao Sun, Erli Meng, Bin Wang and Xipeng Qiu

*Decompose, Fuse and Generate: A Formation-Informed Method for Chinese Definition Generation*

Hua Zheng, Damai Dai, Lei Li, Tianyu Liu, Zhifang Sui, Baobao Chang and Yang Liu

*User-Generated Text Corpus for Evaluating Japanese Morphological Analysis and Lexical Normalization*

Shohei Higashiyama, Masao Utiyama, Taro Watanabe and Eiichiro Sumita

*GPT Perdestry Test: Generating new meanings for new words*

Nikolay Malkin, Sameera Lanka, Pranav Goel, Sudha Rao and Nebojsa Jojic

**18:20–19:40 15E: Sentence-level Semantics and Textual Inference**

*Universal Semantic Tagging for English and Mandarin Chinese*

Wenxi Li, Yiyang Hou, Yajie Ye, Li Liang and Weiwei Sun

*ShadowGNN: Graph Projection Neural Network for Text-to-SQL Parser*

Zhi Chen, Lu Chen, Yanbin Zhao, Ruisheng Cao, Zihan Xu, Su Zhu and Kai Yu

*Contextualized and Generalized Sentence Representations by Contrastive Self-Supervised Learning: A Case Study on Discourse Relation Analysis*

Hirokazu Kiyomaru and Sadao Kurohashi

*AMR Parsing with Action-Pointer Transformer*

Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo and Radu Florian

*NL-EDIT: Correcting Semantic Parse Errors through Natural Language Interaction*

Ahmed Elgohary, Christopher Meek, Matthew Richardson, Adam Fourney, Gonzalo Ramos and Ahmed Hassan Awadallah

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

*Unsupervised Concept Representation Learning for Length-Varying Text Similarity*  
Xuchao Zhang, Bo Zong, Wei Cheng, Jingchao Ni, Yanchi Liu and Haifeng Chen

19:40–21:00 16A: Dialogue and Interactive Systems

*Augmenting Knowledge-grounded Conversations with Sequential Knowledge Transition*

Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao and Yanyan Lan

*Adversarial Self-Supervised Learning for Out-of-Domain Detection*

Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu and Weiran Xu

*Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue State Tracking*

Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho and Rajen Subba

*Hierarchical Transformer for Task Oriented Dialog Systems*

Bishal Santra, Potnuru Anusha and Pawan Goyal

*Measuring the ‘I don’t know’ Problem through the Lens of Gricean Quantity*

Huda Khayrallah and João Sedoc

*[TACL9] Dialogue State Tracking with Incremental Reasoning*

Lizi Liao, Le Hong Long, Yunshan Ma, Wenqiang Lei, Tat-Seng Chua

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

**19:40–21:00 16B: Information Extraction**

*RTFE: A Recursive Temporal Fact Embedding Framework for Temporal Knowledge Graph Completion*

Youri Xu, Haihong E, Meina Song, wenyu song, Xiaodong Lv, wang haotian and yang jinrui

*Open Hierarchical Relation Extraction*

Kai Zhang, Yuan Yao, Ruobing Xie, Xu Han, Zhiyuan Liu, Fen Lin, Leyu Lin and Maosong Sun

*Jointly Extracting Explicit and Implicit Relational Triples with Reasoning Pattern Enhanced Binary Pointer Network*

Yubo Chen, Yunqi Zhang, Changran Hu and Yongfeng Huang

*Multi-Grained Knowledge Distillation for Named Entity Recognition*

Xuan Zhou, Xiao Zhang, Chenyang Tao, Junya Chen, Bing Xu, Wei Wang and Jing Xiao

*SGG: Learning to Select, Guide, and Generate for Keyphrase Generation*

Jing Zhao, Junwei Bao, Yifan Wang, Youzheng Wu, Xiaodong He and Bowen Zhou

*Towards Sentiment and Emotion aided Multi-modal Speech Act Classification in Twitter*

Tulika Saha, Apoorva Upadhyaya, Sriparna Saha and Pushpak Bhattacharyya

**19:40–21:00 16C: Machine Translation**

*Generative Imagination Elevates Machine Translation*

Quanyu Long, Mingxuan Wang and Lei Li

*Non-Autoregressive Translation by Learning Target Categorical Codes*

Yu Bao, Shujian Huang, Tong Xiao, Dongqi Wang, Xinyu Dai and Jiajun CHEN

*Training Data Augmentation for Code-Mixed Translation*

Abhirut Gupta, Aditya Vavre and Sunita Sarawagi

*Rethinking Perturbations in Encoder-Decoders for Fast Training*

Sho Takase and Shun Kiyono

**Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)**

*Context-aware Decoder for Neural Machine Translation using a Target-side Document-Level Language Model*

Amane Sugiyama and Naoki Yoshinaga

*Machine Translated Text Detection Through Text Similarity with Round-Trip Translation*

Hoang-Quoc Nguyen-Son, Tran Thao, Seira Hidano, Ishita Gupta and Shinsaku Kiyomoto

**19:40–21:00 16D: Question Answering**

*TR-BERT: Dynamic Token Reduction for Accelerating BERT Inference*

Deming Ye, Yankai Lin, Yufei Huang and Maosong Sun

*Breadth First Reasoning Graph for Multi-hop Question Answering*

Yongjie Huang and Meng Yang

*Improving Zero-Shot Cross-lingual Transfer for Multilingual Question Answering over Knowledge Graph*

Yucheng Zhou, Xiubo Geng, Tao Shen, Wenqiang Zhang and Daxin Jiang

*RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering*

Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu and Haifeng Wang

*DAGN: Discourse-Aware Graph Network for Logical Reasoning*

Yinya Huang, Meng Fang, Yu Cao, Liwei Wang and Xiaodan Liang

*Designing a Minimal Retrieve-and-Read System for Open-Domain Question Answering*

Sohee Yang and Minjoon Seo

*Unsupervised Multi-hop Question Answering by Question Generation*

Liangming Pan, Wenhua Chen, Wenhan Xiong, Min-Yen Kan and William Yang Wang

Wed 09 Jun 2021 (all times PDT, UTC-7) (continued)

19:40–21:00 16E: Summarization

*Sliding Selector Network with Dynamic Memory for Extractive Summarization of Long Documents*

Peng Cui and Le Hu

*AdaptSum: Towards Low-Resource Domain Adaptation for Abstractive Summarization*

Tiezheng Yu, Zihan Liu and Pascale Fung

*QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization*

Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu and Dragomir Radev

*MM-AVS: A Full-Scale Dataset for Multi-modal Summarization*

Xiyang Fu, Jun Wang and Zhenglu Yang

*MediaSum: A Large-scale Media Interview Dataset for Dialogue Summarization*

Chenguang Zhu, Yang Liu, Jie Mei and Michael Zeng

*Improving Faithfulness in Abstractive Summarization with Contrast Candidate Generation and Selection*

Sihao Chen, Fan Zhang, Kazuo Sone and Dan Roth

*Inference Time Style Control for Summarization*

Shuyang Cao and Lu Wang

