

Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT 2018)

Long Papers

New Orleans, Louisiana, USA
1 – 6 June 2018

Volume 1 of 3

ISBN: 978-1-5108-6346-0

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2018) by the Association for Computational Linguistics
All rights reserved.

Printed by Curran Associates, Inc. (2018)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

<i>Label-Aware Double Transfer Learning for Cross-Specialty Medical Named Entity Recognition</i> Zhenghui Wang, Yanru Qu, Liheng Chen, Jian Shen, Weinan Zhang, Shaodian Zhang, Yimei Gao, Gen Gu, Ken Chen and Yong Yu	1
<i>Neural Fine-Grained Entity Type Classification with Hierarchy-Aware Loss</i> Peng Xu and Denilson Barbosa	16
<i>Joint Bootstrapping Machines for High Confidence Relation Extraction</i> Pankaj Gupta, Benjamin Roth and Hinrich Schütze	26
<i>A Deep Generative Model of Vowel Formant Typology</i> Ryan Cotterell and Jason Eisner	37
<i>Fortification of Neural Morphological Segmentation Models for Polysynthetic Minimal-Resource Lan- guages</i> Katharina Kann, Jesus Manuel Mager Hois, Ivan Vladimir Meza Ruiz and Hinrich Schütze	47
<i>Improving Character-Based Decoding Using Target-Side Morphological Information for Neural Ma- chine Translation</i> Peyman Passban, Qun Liu and Andy Way	58
<i>Parsing Speech: a Neural Approach to Integrating Lexical and Acoustic-Prosodic Information</i> Trang Tran, Shubham Toshniwal, Mohit Bansal, Kevin Gimpel, Karen Livescu and Mari Ostendorf	69
<i>Tied Multitask Learning for Neural Speech Translation</i> Antonios Anastasopoulos and David Chiang	82
<i>Please Clap: Modeling Applause in Campaign Speeches</i> Jon Gillick and David Bamman	92
<i>Attentive Interaction Model: Modeling Changes in View in Argumentation</i> Yohan Jo, Shivani Poddar, Byungsoo Jeon, Qinlan Shen, Carolyn Rose and Graham Neubig . . .	103
<i>Automatic Focus Annotation: Bringing Formal Pragmatics Alive in Analyzing the Information Structure of Authentic Data</i> Ramon Ziai and Detmar Meurers	117
<i>Dear Sir or Madam, May I Introduce the GYAFC Dataset: Corpus, Benchmarks and Metrics for For- mality Style Transfer</i> Sudha Rao and Joel Tetreault	129
<i>Improving Implicit Discourse Relation Classification by Modeling Inter-dependencies of Discourse Units in a Paragraph</i> Zeyu Dai and Ruihong Huang	141
<i>A Deep Ensemble Model with Slot Alignment for Sequence-to-Sequence Natural Language Generation</i> Juraj Juraska, Panagiotis Karagiannis, Kevin Bowden and Marilyn Walker	152
<i>A Melody-Conditioned Lyrics Language Model</i> Kento Watanabe, Yuichiroh Matsubayashi, Satoru Fukayama, Masataka Goto, Kentaro Inui and Tomoyasu Nakano	163

<i>Discourse-Aware Neural Rewards for Coherent Text Generation</i> Antoine Bosselut, Asli Celikyilmaz, Xiaodong He, Jianfeng Gao, Po-Sen Huang and Yejin Choi	173
<i>Natural Answer Generation with Heterogeneous Memory</i> Yao Fu and Yansong Feng	185
<i>Query and Output: Generating Words by Querying Distributed Word Representations for Paraphrase Generation</i> Shuming Ma, Xu Sun, Wei Li, Sujian Li, Wenjie Li and Xuancheng Ren	196
<i>Simplification Using Paraphrases and Context-Based Lexical Substitution</i> Reno Kriz, Eleni Miltsakaki, Marianna Apidianaki and Chris Callison-Burch	207
<i>Zero-Shot Question Generation from Knowledge Graphs for Unseen Predicates and Entity Types</i> Hady Elsahar, Christophe Gravier and Frederique Laforest	218
<i>Automated Essay Scoring in the Presence of Biased Ratings</i> Evelin Amorim, Marcia Cançado and Adriano Veloso	229
<i>Content-Based Citation Recommendation</i> Chandra Bhagavatula, Sergey Feldman, Russell Power and Waleed Ammar	238
<i>Looking Beyond the Surface: A Challenge Set for Reading Comprehension over Multiple Sentences</i> Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay and Dan Roth	252
<i>Neural Automated Essay Scoring and Coherence Modeling for Adversarially Crafted Input</i> Younna Farag, Helen Yannakoudakis and Ted Briscoe	263
<i>QuickEdit: Editing Text & Translations by Crossing Words Out</i> David Grangier and Michael Auli	272
<i>Tempo-Lexical Context Driven Word Embedding for Cross-Session Search Task Extraction</i> Procheta Sen, Debasis Ganguly and Gareth Jones	283
<i>Zero-Shot Sequence Labeling: Transferring Knowledge from Sentences to Tokens</i> Marek Rei and Anders Søgaard	293
<i>Variable Typing: Assigning Meaning to Variables in Mathematical Text</i> Yiannos Stathopoulos, Simon Baker, Marek Rei and Simone Teufel	303
<i>Learning beyond Datasets: Knowledge Graph Augmented Neural Networks for Natural Language Processing</i> Annervaz K M, Somnath Basu Roy Chowdhury and Ambedkar Dukkipati	313
<i>Comparing Constraints for Taxonomic Organization</i> Anne Cocos, Marianna Apidianaki and Chris Callison-Burch	323
<i>Improving Lexical Choice in Neural Machine Translation</i> Toan Nguyen and David Chiang	334
<i>Universal Neural Machine Translation for Extremely Low Resource Languages</i> Jiatao Gu, Hany Hassan, Jacob Devlin and Victor O.K. Li	344
<i>Classical Structured Prediction Losses for Sequence to Sequence Learning</i> Sergey Edunov, Myle Ott, Michael Auli, David Grangier and Marc’Aurelio Ranzato	355

<i>Deep Dirichlet Multinomial Regression</i> Adrian Benton and Mark Dredze	365
<i>Microblog Conversation Recommendation via Joint Modeling of Topics and Discourse</i> Kingshan Zeng, Jing Li, Lu Wang, Nicholas Beauchamp, Sarah Shugars and Kam-Fai Wong ..	375
<i>Before Name-Calling: Dynamics and Triggers of Ad Hominem Fallacies in Web Argumentation</i> Ivan Habernal, Henning Wachsmuth, Iryna Gurevych and Benno Stein	386
<i>Scene Graph Parsing as Dependency Parsing</i> Yu-Siang Wang, Chenxi Liu, Xiaohui Zeng and Alan Yuille	397
<i>Learning Visually Grounded Sentence Representations</i> Douwe Kiela, Alexis Conneau, Allan Jabri and Maximilian Nickel	408
<i>Comparatives, Quantifiers, Proportions: a Multi-Task Model for the Learning of Quantities from Vision</i> Sandro Pezzelle, Ionut-Teodor Sorodoc and Raffaella Bernardi	419
<i>Being Negative but Constructively: Lessons Learnt from Creating Better Visual Question Answering Datasets</i> Wei-Lun Chao, Hexiang Hu and Fei Sha	431
<i>Abstract Meaning Representation for Paraphrase Detection</i> Fuad Issa, Marco Damonte, Shay B. Cohen, Xiaohui Yan and Yi Chang	442
<i>attr2vec: Jointly Learning Word and Contextual Attribute Embeddings with Factorization Machines</i> Fabio Petroni, Vassilis Plachouras, Timothy Nugent and Jochen L. Leidner	453
<i>Can Network Embedding of Distributional Thesaurus Be Combined with Word Vectors for Better Representation?</i> Abhik Jana and Pawan Goyal	463
<i>Deep Neural Models of Semantic Shift</i> Alex Rosenfeld and Katrin Erk	474
<i>Distributional Inclusion Vector Embedding for Unsupervised Hypernymy Detection</i> Haw-Shiuan Chang, Ziyun Wang, Luke Vilnis and Andrew McCallum	485
<i>Mining Possessions: Existence, Type and Temporal Anchors</i> Dhivya Chinnappa and Eduardo Blanco	496
<i>Neural Tensor Networks with Diagonal Slice Matrices</i> Takahiro Ishihara, Katsuhiko Hayashi, Hitoshi Manabe, Masashi Shimbo and Masaaki Nagata	506
<i>Post-Specialisation: Retrofitting Vectors of Words Unseen in Lexical Resources</i> Ivan Vulić, Goran Glavaš, Nikola Mrkšić and Anna Korhonen	516
<i>Unsupervised Learning of Sentence Embeddings Using Compositional n-Gram Features</i> Matteo Pagliardini, Prakhar Gupta and Martin Jaggi	528
<i>Learning Domain Representation for Multi-Domain Sentiment Classification</i> Qi Liu, Yue Zhang and Jiangming Liu	541
<i>Learning Sentence Representations over Tree Structures for Target-Dependent Classification</i> Junwen Duan, Xiao Ding and Ting Liu	551

<i>Relevant Emotion Ranking from Text Constrained with Emotion Relationships</i> Deyu Zhou, Yang Yang and Yulan He	561
<i>Solving Data Sparsity for Aspect Based Sentiment Analysis Using Cross-Linguality and Multi-Linguality</i> Md Shad Akhtar, Palaash Sawant, Sukanta Sen, Asif Ekbal and Pushpak Bhattacharyya	572
<i>SRL4ORL: Improving Opinion Role Labeling Using Multi-Task Learning with Semantic Role Labeling</i> Ana Marasović and Anette Frank	583
<i>Approaching Neural Grammatical Error Correction as a Low-Resource Machine Translation Task</i> Marcin Junczys-Dowmunt, Roman Grundkiewicz, Shubha Guha and Kenneth Heafield	595
<i>Robust Cross-Lingual Hypernymy Detection Using Dependency Context</i> Shyam Upadhyay, Yogarshi Vyas, Marine Carpuat and Dan Roth	607
<i>Noising and Denoising Natural Language: Diverse Backtranslation for Grammar Correction</i> Ziang Xie, Guillaume Genthial, Stanley Xie, Andrew Ng and Dan Jurafsky	619
<i>Self-Training for Jointly Learning to Ask and Answer Questions</i> Mrinmaya Sachan and Eric Xing	629
<i>The Web as a Knowledge-Base for Answering Complex Questions</i> Alon Talmor and Jonathan Berant	641
<i>A Meaning-Based Statistical English Math Word Problem Solver</i> Chao-Chun Liang, Yu-Shiang Wong, Yi-Chung Lin and Keh-Yih Su	652
<i>Fine-Grained Temporal Orientation and its Relationship with Psycho-Demographic Correlates</i> Sabyasachi Kamila, Mohammed Hasanuzzaman, Asif Ekbal, Pushpak Bhattacharyya and Andy Way	663
<i>Querying Word Embeddings for Similarity and Relatedness</i> Fatemeh Torabi Asr, Robert Zinkov and Michael Jones	675
<i>Semantic Structural Evaluation for Text Simplification</i> Elior Sulem, Omri Abend and Ari Rappoport	685
<i>Entity Commonsense Representation for Neural Abstractive Summarization</i> Reinald Kim Amplayo, Seonjae Lim and Seung-won Hwang	697
<i>Newsroom: A Dataset of 1.3 Million Summaries with Diverse Extractive Strategies</i> Max Grusky, Mor Naaman and Yoav Artzi	708
<i>Polyglot Semantic Parsing in APIs</i> Kyle Richardson, Jonathan Berant and Jonas Kuhn	720
<i>Neural Models of Factuality</i> Rachel Rudinger, Aaron Steven White and Benjamin Van Durme	731
<i>Accurate Text-Enhanced Knowledge Graph Representation Learning</i> Bo An, Bo Chen, Xianpei Han and Le Sun	745
<i>Acquisition of Phrase Correspondences Using Natural Deduction Proofs</i> Hitomi Yanaka, Koji Mineshima, Pascual Martínez-Gómez and Daisuke Bekki	756

<i>Automatic Stance Detection Using End-to-End Memory Networks</i> Mitra Mohtarami, Ramy Baly, James Glass, Preslav Nakov, Lluís Màrquez and Alessandro Moschitti	767
<i>Collective Entity Disambiguation with Structured Gradient Tree Boosting</i> Yi Yang, Ozan Irsoy and Kazi Shefaet Rahman	777
<i>DeepAlignment: Unsupervised Ontology Matching with Refined Word Vectors</i> Prodromos Kolyvakis, Alexandros Kalousis and Dimitris Kiritsis	787
<i>Efficient Sequence Learning with Group Recurrent Networks</i> Fei Gao, Lijun Wu, Li Zhao, Tao Qin, Xueqi Cheng and Tie-Yan Liu	799
<i>FEVER: a Large-scale Dataset for Fact Extraction and VERification</i> James Thorne, Andreas Vlachos, Christos Christodoulopoulos and Arpit Mittal	809
<i>Global Relation Embedding for Relation Extraction</i> Yu Su, Honglei Liu, Semih Yavuz, Izzeddin Gur, Huan Sun and Xifeng Yan	820
<i>Implicit Argument Prediction with Event Knowledge</i> Pengxiang Cheng and Katrin Erk	831
<i>Improving Temporal Relation Extraction with a Globally Acquired Statistical Resource</i> Qiang Ning, Hao Wu, Haoruo Peng and Dan Roth	841
<i>Multimodal Named Entity Recognition for Short Social Media Posts</i> Seungwhan Moon, Leonardo Neves and Vitor Carvalho	852
<i>Nested Named Entity Recognition Revisited</i> Arzoo Katiyar and Claire Cardie	861
<i>Simultaneously Self-Attending to All Mentions for Full-Abstract Biological Relation Extraction</i> Patrick Verga, Emma Strubell and Andrew McCallum	872
<i>Supervised Open Information Extraction</i> Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer and Ido Dagan	885
<i>Embedding Syntax and Semantics of Prepositions via Tensor Decomposition</i> Hongyu Gong, Suma Bhat and Pramod Viswanath	896
<i>From Phonology to Syntax: Unsupervised Linguistic Typology at Different Levels with Language Embeddings</i> Johannes Bjerva and Isabelle Augenstein	907
<i>Monte Carlo Syntax Marginals for Exploring and Using Dependency Parses</i> Katherine Keith, Su Lin Blodgett and Brendan O'Connor	917
<i>Neural Particle Smoothing for Sampling from Conditional Sequence Models</i> Chu-Cheng Lin and Jason Eisner	929
<i>Neural Syntactic Generative Models with Exact Marginalization</i> Jan Buys and Phil Blunsom	942
<i>Noise-Robust Morphological Disambiguation for Dialectal Arabic</i> Nasser Zalmout, Alexander Erdmann and Nizar Habash	953

<i>Parsing Tweets into Universal Dependencies</i> Yijia Liu, Yi Zhu, Wanxiang Che, Bing Qin, Nathan Schneider and Noah A. Smith	965
<i>Robust Multilingual Part-of-Speech Tagging via Adversarial Training</i> Michihiro Yasunaga, Jungo Kasai and Dragomir Radev	976
<i>Universal Dependency Parsing for Hindi-English Code-Switching</i> Irshad Bhat, Riyaz A. Bhat, Manish Shrivastava and Dipti Sharma	987
<i>What's Going On in Neural Constituency Parsers? An Analysis</i> David Gaddy, Mitchell Stern and Dan Klein	999
<i>Deep Generative Model for Joint Alignment and Word Representation</i> Miguel Rios, Wilker Aziz and Khalil Simaan	1011
<i>Learning Word Embeddings for Low-Resource Languages by PU Learning</i> Chao Jiang, Hsiang-Fu Yu, Cho-Jui Hsieh and Kai-Wei Chang	1024
<i>Exploring the Role of Prior Beliefs for Argument Persuasion</i> Esin Durmus and Claire Cardie	1035
<i>Inducing a Lexicon of Abusive Words – a Feature-Based Approach</i> Michael Wiegand, Josef Ruppenhofer, Anna Schmidt and Clayton Greenberg	1046
<i>Author Commitment and Social Power: Automatic Belief Tagging to Infer the Social Context of Interactions</i> Vinodkumar Prabhakaran, Premkumar Ganeshkumar and Owen Rambow	1057
<i>Comparing Automatic and Human Evaluation of Local Explanations for Text Classification</i> Dong Nguyen	1069
<i>Deep Temporal-Recurrent-Replicated-Softmax for Topical Trends over Time</i> Pankaj Gupta, Subburam Rajaram, Hinrich Schütze and Bernt Andrassy	1079
<i>Lessons from the Bible on Modern Topics: Low-Resource Multilingual Topic Model Evaluation</i> Shudong Hao, Jordan Boyd-Graber and Michael J. Paul	1090
<i>Explainable Prediction of Medical Codes from Clinical Text</i> James Mullenbach, Sarah Wiegreffe, Jon Duke, Jimeng Sun and Jacob Eisenstein	1101
<i>A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference</i> Adina Williams, Nikita Nangia and Samuel Bowman	1112
<i>Filling Missing Paths: Modeling Co-occurrences of Word Pairs and Dependency Paths for Recognizing Lexical Semantic Relations</i> Koki Washio and Tsuneaki Kato	1123
<i>Specialising Word Vectors for Lexical Entailment</i> Ivan Vulić and Nikola Mrkšić	1134
<i>Cross-Lingual Abstract Meaning Representation Parsing</i> Marco Damonte and Shay B. Cohen	1146
<i>Sentences with Gapping: Parsing and Reconstructing Elided Predicates</i> Sebastian Schuster, Joakim Nivre and Christopher D. Manning	1156

<i>A Structured Syntax-Semantics Interface for English-AMR Alignment</i> Ida Szubert, Adam Lopez and Nathan Schneider	1169
<i>End-to-End Graph-Based TAG Parsing with Neural Networks</i> Jungo Kasai, Robert Frank, Pauli Xu, William Merrill and Owen Rambow	1181
<i>Colorless Green Recurrent Networks Dream Hierarchically</i> Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen and Marco Baroni	1195
<i>Diverse Few-Shot Text Classification with Multiple Metrics</i> Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang and Bowen Zhou	1206
<i>Early Text Classification Using Multi-Resolution Concept Representations</i> Adrian Pastor López Monroy, Fabio A. González, Manuel Montes, Hugo Jair Escalante and Tamar Solorio	1216
<i>Multinomial Adversarial Networks for Multi-Domain Text Classification</i> Xilun Chen and Claire Cardie	1226
<i>Pivot Based Language Modeling for Improved Neural Domain Adaptation</i> Yftah Ziser and Roi Reichart	1241
<i>Reinforced Co-Training</i> Jiawei Wu, Lei Li and William Yang Wang	1252
<i>Tensor Product Generation Networks for Deep NLP Modeling</i> Qiuyuan Huang, Paul Smolensky, Xiaodong He, Li Deng and Dapeng Wu	1263
<i>The Context-Dependent Additive Recurrent Neural Net</i> Quan Hung Tran, Tuan Lai, Gholamreza Haffari, Ingrid Zukerman, Trung Bui and Hung Bui	1274
<i>Combining Character and Word Information in Neural Machine Translation Using a Multi-Level Attention</i> Huadong Chen, Shujian Huang, David Chiang, Xinyu Dai and Jiajun Chen	1284
<i>Dense Information Flow for Neural Machine Translation</i> Yanyao Shen, Xu Tan, Di He, Tao Qin and Tie-Yan Liu	1294
<i>Evaluating Discourse Phenomena in Neural Machine Translation</i> Rachel Bawden, Rico Sennrich, Alexandra Birch and Barry Haddow	1304
<i>Fast Lexically Constrained Decoding with Dynamic Beam Allocation for Neural Machine Translation</i> Matt Post and David Vilar	1314
<i>Guiding Neural Machine Translation with Retrieved Translation Pieces</i> Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig and Satoshi Nakamura	1325
<i>Handling Homographs in Neural Machine Translation</i> Frederick Liu, Han Lu and Graham Neubig	1336
<i>Improving Neural Machine Translation with Conditional Sequence Generative Adversarial Nets</i> Zhen Yang, Wei Chen, Feng Wang and Bo Xu	1346
<i>Neural Machine Translation for Bilingually Scarce Scenarios: a Deep Multi-Task Learning Approach</i> Poorya Zareemoodi and Gholamreza Haffari	1356

<i>Self-Attentive Residual Decoder for Neural Machine Translation</i>	
Lesly Miculicich Werlen, Nikolaos Pappas, Dhananjay Ram and Andrei Popescu-Belis	1366
<i>Target Foresight Based Attention for Neural Machine Translation</i>	
Xintong Li, Lemao Liu, Zhaopeng Tu, Shuming Shi and Max Meng	1380
<i>Context Sensitive Neural Lemmatization with Lematus</i>	
Toms Bergmanis and Sharon Goldwater	1391
<i>Modeling Noisiness to Recognize Named Entities using Multitask Neural Networks on Social Media</i>	
Gustavo Aguilar, Adrian Pastor López Monroy, Fabio González and Thamar Solorio	1401
<i>Reusing Weights in Subword-Aware Neural Language Models</i>	
Zhenisbek Assylbekov and Rustem Takhanov	1413
<i>Simple Models for Word Formation in Slang</i>	
Vivek Kulkarni and William Yang Wang	1424
<i>Using Morphological Knowledge in Open-Vocabulary Neural Language Models</i>	
Austin Matthews, Graham Neubig and Chris Dyer	1435
<i>A Neural Layered Model for Nested Named Entity Recognition</i>	
Meizhi Ju, Makoto Miwa and Sophia Ananiadou	1446
<i>DR-BiLSTM: Dependent Reading Bidirectional LSTM for Natural Language Inference</i>	
Reza Ghaeini, Sadid A. Hasan, Vivek Datla, Joey Liu, Kathy Lee, Ashequl Qadir, Yuan Ling, Aaditya Prakash, Xiaoli Fern and Oladimeji Farri	1460
<i>KBGAN: Adversarial Learning for Knowledge Graph Embeddings</i>	
Liwei Cai and William Yang Wang	1470
<i>Multimodal Frame Identification with Multilingual Evaluation</i>	
Teresa Botschen, Iryna Gurevych, Jan-Christoph Klie, Hatem Mousselly Sergieh and Stefan Roth	1481
<i>Learning Joint Semantic Parsers from Disjoint Data</i>	
Hao Peng, Sam Thomson, Swabha Swayamdipta and Noah A. Smith	1492
<i>Identifying Semantic Divergences in Parallel Text without Annotations</i>	
Yogarshi Vyas, Xing Niu and Marine Carpuat	1503
<i>Bootstrapping Generators from Noisy Data</i>	
Laura Perez-Beltrachini and Mirella Lapata	1516
<i>SHAPED: Shared-Private Encoder-Decoder for Text Style Adaptation</i>	
Ye Zhang, Nan Ding and Radu Soricut	1528
<i>Generating Descriptions from Structured Data Using a Bifocal Attention Mechanism and Gated Orthogonalization</i>	
Preksha Nema, Shreyas Shetty, Parag Jain, Anirban Laha, Karthik Sankaranarayanan and Mitesh M. Khapra	1539
<i>CliCR: a Dataset of Clinical Case Reports for Machine Reading Comprehension</i>	
Simon Suster and Walter Daelemans	1551

<i>Learning to Collaborate for Question Answering and Asking</i> Duyu Tang, Nan Duan, Zhao Yan, Zhirui Zhang, Yibo Sun, Shujie Liu, Yuanhua Lv and Ming Zhou	1564
<i>Learning to Rank Question-Answer Pairs Using Hierarchical Recurrent Encoder with Latent Topic Clustering</i> Seunghyun Yoon, Joongbo Shin and Kyomin Jung	1575
<i>Supervised and Unsupervised Transfer Learning for Question Answering</i> Yu-An Chung, Hung-yi Lee and James Glass	1585
<i>Tracking State Changes in Procedural Text: a Challenge Dataset and Models for Process Paragraph Comprehension</i> Bhavana Dalvi, Lifu Huang, Niket Tandon, Wen-tau Yih and Peter Clark	1595
<i>Combining Deep Learning and Topic Modeling for Review Understanding in Context-Aware Recommendation</i> Mingmin Jin, Xin Luo, Huiling Zhu and Hankz Hankui Zhuo	1605
<i>Deconfounded Lexicon Induction for Interpretable Social Science</i> Reid Pryzant, Kelly Shen, Dan Jurafsky and Stefan Wagner	1615
<i>Detecting Denial-of-Service Attacks from Social Media Text: Applying NLP to Computer Security</i> Nathanael Chambers, Ben Fry and James McMasters	1626
<i>The Importance of Calibration for Estimating Proportions from Annotations</i> Dallas Card and Noah A. Smith	1636
<i>A Dataset of Peer Reviews (PeerRead): Collection, Insights and NLP Applications</i> Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy and Roy Schwartz	1647
<i>Deep Communicating Agents for Abstractive Summarization</i> Asli Celikyilmaz, Antoine Bosselut, Xiaodong He and Yejin Choi	1662
<i>Encoding Conversation Context for Neural Keyphrase Extraction from Microblog Posts</i> Yingyi Zhang, Jing Li, Yan Song and Chengzhi Zhang	1676
<i>Estimating Summary Quality with Pairwise Preferences</i> Markus Zopf	1687
<i>Generating Topic-Oriented Summaries Using Neural Attention</i> Kundan Krishna and Balaji Vasani Srinivasan	1697
<i>Generative Bridging Network for Neural Sequence Prediction</i> Wenhu Chen, Guanlin Li, Shuo Ren, Shujie Liu, Zhirui Zhang, Mu Li and Ming Zhou	1706
<i>Higher-Order Syntactic Attention Network for Longer Sentence Compression</i> Hidetaka Kamigaito, Katsuhiko Hayashi, Tsutomu Hirao and Masaaki Nagata	1716
<i>Neural Storyline Extraction Model for Storyline Generation from News Articles</i> Deyu Zhou, Linsen Guo and Yulan He	1727
<i>Provable Fast Greedy Compressive Summarization with Any Monotone Submodular Function</i> Shinsaku Sakaue, Tsutomu Hirao, Masaaki Nishino and Masaaki Nagata	1737

<i>Ranking Sentences for Extractive Summarization with Reinforcement Learning</i> Shashi Narayan, Shay B. Cohen and Mirella Lapata	1747
<i>Relational Summarization for Corpus Analysis</i> Abram Handler and Brendan O'Connor	1760
<i>What's This Movie About? A Joint Neural Network Architecture for Movie Content Analysis</i> Philip John Gorinski and Mirella Lapata	1770
<i>Which Scores to Predict in Sentence Regression for Text Summarization?</i> Markus Zopf, Eneldo Loza Mencía and Johannes Fürnkranz	1782
<i>A Hierarchical Latent Structure for Variational Conversation Modeling</i> Yookoon Park, Jaemin Cho and Gunhee Kim	1792
<i>Detecting Egregious Conversations between Customers and Virtual Agents</i> Tommy Sandbank, Michal Shmueli-Scheuer, Jonathan Herzig, David Konopnicki, John Richards and David Piorkowski	1802
<i>Learning to Disentangle Interleaved Conversational Threads with a Siamese Hierarchical Network and Similarity Ranking</i> Jyun-Yu Jiang, Francine Chen, Yan-Ying Chen and Wei Wang	1812
<i>Variational Knowledge Graph Reasoning</i> Wenhu Chen, Wenhan Xiong, Xifeng Yan and William Yang Wang	1823
<i>Inducing Temporal Relations from Time Anchor Annotation</i> Fei Cheng and Yusuke Miyao	1833
<i>ELDEN: Improved Entity Linking Using Densified Knowledge Graphs</i> Priya Radhakrishnan, Partha Talukdar and Vasudeva Varma	1844
<i>Interpretable Charge Predictions for Criminal Cases: Learning to Generate Court Views from Fact Descriptions</i> Hai Ye, Xin Jiang, Zhunchen Luo and Wenhan Chao	1854
<i>Delete, Retrieve, Generate: a Simple Approach to Sentiment and Style Transfer</i> Juncen Li, Robin Jia, He He and Percy Liang	1865
<i>Adversarial Example Generation with Syntactically Controlled Paraphrase Networks</i> Mohit Iyyer, John Wieting, Kevin Gimpel and Luke Zettlemoyer	1875
<i>Sentiment Analysis: It's Complicated!</i> Kian Kenyon-Dean, Eisha Ahmed, Scott Fujimoto, Jeremy Georges-Filteau, Christopher Glasz, Barleen Kaur, Auguste Lalande, Shruti Bhanderi, Robert Belfer, Nirmal Kanagasabai, Roman Sarrazin- gendron, Rohit Verma and Derek Ruths	1886
<i>Multi-Task Learning of Pairwise Sequence Classification Tasks over Disparate Label Spaces</i> Isabelle Augenstein, Sebastian Ruder and Anders Søgaard	1896
<i>Word Emotion Induction for Multiple Languages as a Deep Multi-Task Learning Problem</i> Sven Buechel and Udo Hahn	1907
<i>Human Needs Categorization of Affective Events Using Labeled and Unlabeled Data</i> Haibo Ding and Ellen Riloff	1919

<i>The Argument Reasoning Comprehension Task: Identification and Reconstruction of Implicit Warrants</i> Ivan Habernal, Henning Wachsmuth, Iryna Gurevych and Benno Stein	1930
<i>Linguistic Cues to Deception and Perceived Deception in Interview Dialogues</i> Sarah Ita Levitan, Angel Maredia and Julia Hirschberg	1941
<i>Unified Pragmatic Models for Generating and Following Instructions</i> Daniel Fried, Jacob Andreas and Dan Klein	1951
<i>Hierarchical Structured Model for Fine-to-Coarse Manifesto Text Analysis</i> Shivashankar Subramanian, Trevor Cohn and Timothy Baldwin	1964
<i>Behavior Analysis of NLI Models: Uncovering the Influence of Three Factors on Robustness</i> Ivan Sanchez, Jeff Mitchell and Sebastian Riedel	1975
<i>Assessing Language Proficiency from Eye Movements in Reading</i> Yevgeni Berzak, Boris Katz and Roger Levy	1986
<i>Comparing Theories of Speaker Choice Using a Model of Classifier Production in Mandarin Chinese</i> Meilin Zhan and Roger Levy	1997
<i>Spotting Spurious Data with Neural Networks</i> Hadi Amiri, Timothy Miller and Guergana Savova	2006
<i>The Timing of Lexical Memory Retrievals in Language Production</i> Jeremy Cole and David Reitter	2017
<i>Unsupervised Induction of Linguistic Categories with Records of Reading, Speaking, and Writing</i> Maria Barrett, Ana Valeria Gonzalez-Garduño, Lea Frermann and Anders Søgaard	2028
<i>Challenging Reading Comprehension on Daily Conversation: Passage Completion on Multiparty Dialog</i> Kaixin Ma, Tomasz Jurczyk and Jinho D. Choi	2039
<i>Dialog Generation Using Multi-Turn Reasoning Neural Networks</i> Xianchao Wu, Ander Martinez and Momo Klyen	2049
<i>Dialogue Learning with Human Teaching and Feedback in End-to-End Trainable Task-Oriented Dialogue Systems</i> Bing Liu, Gokhan Tür, Dilek Hakkani-Tür, Pararth Shah and Larry Heck	2060
<i>LSDSCC: a Large Scale Domain-Specific Conversational Corpus for Response Generation with Diversity Oriented Evaluation Metrics</i> Zhen Xu, Nan Jiang, Bingquan Liu, Wenge Rong, Bowen Wu, Baoxun Wang, Zhuoran Wang and Xiaolong Wang	2070
<i>EMR Coding with Semi-Parametric Multi-Head Matching Networks</i> Anthony Rios and Ramakanth Kavuluru	2081
<i>Factors Influencing the Surprising Instability of Word Embeddings</i> Laura Wendlandt, Jonathan K. Kummerfeld and Rada Mihalcea	2092
<i>Mining Evidences for Concept Stock Recommendation</i> Qi Liu and Yue Zhang	2103
<i>Binarized LSTM Language Model</i> Xuan Liu, Di Cao and Kai Yu	2113

<i>Conversational Memory Network for Emotion Recognition in Dyadic Dialogue Videos</i> Devamanyu Hazarika, Soujanya Poria, Amir Zadeh, Erik Cambria, Louis-Philippe Morency and Roger Zimmermann	2122
<i>How Time Matters: Learning Time-Decay Attention for Contextual Spoken Language Understanding in Dialogues</i> Shang-Yu Su, Pei-Chieh Yuan and Yun-Nung Chen	2133
<i>Towards Understanding Text Factors in Oral Reading</i> Anastassia Loukina, Van Rynald T. Licalalde and Beata Beigman Klebanov	2143
<i>Generating Bilingual Pragmatic Color References</i> Will Monroe, Jennifer Hu, Andrew Jong and Christopher Potts	2155
<i>Learning with Latent Language</i> Jacob Andreas, Dan Klein and Sergey Levine	2166
<i>Object Counts! Bringing Explicit Detections Back into Image Captioning</i> Josiah Wang, Pranava Swaroop Madhyastha and Lucia Specia	2180
<i>Quantifying the Visual Concreteness of Words and Topics in Multimodal Datasets</i> Jack Hessel, David Mimno and Lillian Lee	2194
<i>Speaker Naming in Movies</i> Mahmoud Azab, Mingzhe Wang, Max Smith, Noriyuki Kojima, Jia Deng and Rada Mihalcea	2206
<i>Stacking with Auxiliary Features for Visual Question Answering</i> Nazneen Fatema Rajani and Raymond Mooney	2217
<i>Deep Contextualized Word Representations</i> Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee and Luke Zettlemoyer	2227
<i>Learning to Map Context-Dependent Sentences to Executable Formal Queries</i> Alane Suhr, Srinivasan Iyer and Yoav Artzi	2238
<i>Neural Text Generation in Stories Using Entity Representations as Context</i> Elizabeth Clark, Yangfeng Ji and Noah A. Smith	2250
<i>Recurrent Neural Networks as Weighted Language Recognizers</i> Yining Chen, Sorcha Gilroy, Andreas Maletti, Jonathan May and Kevin Knight	2261