

55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)

Long Papers

Vancouver, Canada
30 July – 4 August 2017

Volume 1 of 3

ISBN: 978-1-5108-4567-1

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© (2017) by the Association for Computational Linguistics
All rights reserved.

Printed by Curran Associates, Inc. (2017)

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>Adversarial Multi-task Learning for Text Classification</i>	
Pengfei Liu, Xipeng Qiu and Xuanjing Huang	1
<i>Neural End-to-End Learning for Computational Argumentation Mining</i>	
Steffen Eger, Johannes Daxenberger and Iryna Gurevych	11
<i>Neural Symbolic Machines: Learning Semantic Parsers on Freebase with Weak Supervision</i>	
Chen Liang, Jonathan Berant, Quoc Le, Kenneth D. Forbus and Ni Lao	23
<i>Neural Relation Extraction with Multi-lingual Attention</i>	
Yankai Lin, Zhiyuan Liu and Maosong Sun	34
<i>Learning Structured Natural Language Representations for Semantic Parsing</i>	
Jianpeng Cheng, Siva Reddy, Vijay Saraswat and Mirella Lapata	44
<i>Morph-fitting: Fine-Tuning Word Vector Spaces with Simple Language-Specific Rules</i>	56
Ivan Vulić, Nikola Mrkšić, Roi Reichart, Diarmuid Ó Séaghdha, Steve Young and Anna Korhonen	
<i>Skip-Gram - Zipf + Uniform = Vector Additivity</i>	
Alex Gittens, Dimitris Achlioptas and Michael W. Mahoney	69
<i>The State of the Art in Semantic Representation</i>	
Omri Abend and Ari Rappoport	77
<i>Joint Learning for Event Coreference Resolution</i>	
Jing Lu and Vincent Ng	90
<i>Generating and Exploiting Large-scale Pseudo Training Data for Zero Pronoun Resolution</i>	
Ting Liu, Yiming Cui, Qingyu Yin, Wei-Nan Zhang, Shijin Wang and Guoping Hu	102
<i>Discourse Mode Identification in Essays</i>	
Wei Song, Dong Wang, Ruiji Fu, Lizhen Liu, Ting Liu and Guoping Hu	112
<i>A Convolutional Encoder Model for Neural Machine Translation</i>	
Jonas Gehring, Michael Auli, David Grangier and Yann Dauphin	123
<i>Deep Neural Machine Translation with Linear Associative Unit</i>	
Mingxuan Wang, Zhengdong Lu, Jie Zhou and Qun Liu	136
<i>Neural AMR: Sequence-to-Sequence Models for Parsing and Generation</i>	
Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi and Luke Zettlemoyer	146
<i>Program Induction by Rationale Generation: Learning to Solve and Explain Algebraic Word Problems</i>	
Wang Ling, Dani Yogatama, Chris Dyer and Phil Blunsom	158
<i>Automatically Generating Rhythmic Verse with Neural Networks</i>	
Jack Hopkins and Douwe Kiela	168
<i>Creating Training Corpora for NLG Micro-Planners</i>	
Claire Gardent, Anastasia Shimorina, Shashi Narayan and Laura Perez-Beltrachini	179

<i>Gated Self-Matching Networks for Reading Comprehension and Question Answering</i> Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang and Ming Zhou	189
<i>Generating Natural Answers by Incorporating Copying and Retrieving Mechanisms in Sequence-to-Sequence Learning</i> Shizhu He, Cao Liu, Kang Liu and Jun Zhao	199
<i>Coarse-to-Fine Question Answering for Long Documents</i> Eunsol Choi, Daniel Hewlett, Jakob Uszkoreit, Illia Polosukhin, Alexandre Lacoste and Jonathan Berant	209
<i>An End-to-End Model for Question Answering over Knowledge Base with Cross-Attention Combining Global Knowledge</i> Yanchao Hao, Yuanzhe Zhang, Kang Liu, Shizhu He, Zhanyi Liu, Hua Wu and Jun Zhao	221
<i>Translating NeuraLeSe</i> Jacob Andreas, Anca Dragan and Dan Klein	232
<i>Obtaining referential word meanings from visual and distributional information: Experiments on object naming</i> Sina Zarrieß and David Schlangen	243
<i>FOIL it! Find One mismatch between Image and Language caption</i> Ravi Shekhar, Sandro Pezzelle, Yauhen Klimovich, Aurélie Herbelot, Moin Nabi, Enver Sangineto and Raffaella Bernardi	255
<i>Verb Physics: Relative Physical Knowledge of Actions and Objects</i> Maxwell Forbes and Yejin Choi	266
<i>A* CCG Parsing with a Supertag and Dependency Factored Model</i> Masashi Yoshikawa, Hiroshi Noji and Yuji Matsumoto	277
<i>A Full Non-Monotonic Transition System for Unrestricted Non-Projective Parsing</i> Daniel Fernández-González and Carlos Gómez-Rodríguez	288
<i>Aggregating and Predicting Sequence Labels from Crowd Annotations</i> An Thanh Nguyen, Byron Wallace, Junyi Jessy Li, Ani Nenkova and Matthew Lease	299
<i>Multi-space Variational Encoder-Decoders for Semi-supervised Labeled Sequence Transduction</i> Chunting Zhou and Graham Neubig	310
<i>Scalable Bayesian Learning of Recurrent Neural Networks for Language Modeling</i> Zhe Gan, Chunyuan Li, Changyou Chen, Yunchen Pu, Qinliang Su and Lawrence Carin	321
<i>Learning attention for historical text normalization by learning to pronounce</i> Marcel Bollmann, Joachim Bingel and Anders Søgaard	332
<i>Deep Learning in Semantic Kernel Spaces</i> Danilo Croce, Simone Filice, Giuseppe Castellucci and Roberto Basili	345
<i>Topically Driven Neural Language Model</i> Jey Han Lau, Timothy Baldwin and Trevor Cohn	355
<i>Handling Cold-Start Problem in Review Spam Detection by Jointly Embedding Texts and Behaviors</i> Xuepeng Wang, Kang Liu and Jun Zhao	366

<i>Learning Cognitive Features from Gaze Data for Sentiment and Sarcasm Classification using Convolutional Neural Network</i>	
Abhijit Mishra, Kuntal Dey and Pushpak Bhattacharyya	377
<i>An Unsupervised Neural Attention Model for Aspect Extraction</i>	
Ruidan He, Wee Sun Lee, Hwee Tou Ng and Daniel Dahlmeier	388
<i>Other Topics You May Also Agree or Disagree: Modeling Inter-Topic Preferences using Tweets and Matrix Factorization</i>	
Akira Sasaki, Kazuaki Hanawa, Naoaki Okazaki and Kentaro Inui	398
<i>Automatically Labeled Data Generation for Large Scale Event Extraction</i>	
Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu and Jun Zhao	409
<i>Time Expression Analysis and Recognition Using Syntactic Token Types and General Heuristic Rules</i>	
Xiaoshi Zhong, Aixin Sun and Erik Cambria	420
<i>Learning with Noise: Enhance Distantly Supervised Relation Extraction with Dynamic Transition Matrix</i>	
Bingfeng Luo, Yansong Feng, Zheng Wang, Zhanxing Zhu, Songfang Huang, Rui Yan and Dongyan Zhao	430
<i>A Syntactic Neural Model for General-Purpose Code Generation</i>	
Pengcheng Yin and Graham Neubig	440
<i>Learning bilingual word embeddings with (almost) no bilingual data</i>	
Mikel Artetxe, Gorka Labaka and Eneko Agirre	451
<i>Abstract Meaning Representation Parsing using LSTM Recurrent Neural Networks</i>	
William Foland and James H. Martin	463
<i>Deep Semantic Role Labeling: What Works and What's Next</i>	
Luheng He, Kenton Lee, Mike Lewis and Luke Zettlemoyer	473
<i>Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access</i>	
Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed and Li Deng	484
<i>Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-Based Chatbots</i>	
Yu Wu, Wei Wu, Chen Xing, Ming Zhou and Zhoujun Li	496
<i>Learning Word-Like Units from Joint Audio-Visual Analysis</i>	
David Harwath and James Glass	506
<i>Joint CTC/attention decoding for end-to-end speech recognition</i>	
Takaaki Hori, Shinji Watanabe and John Hershey	518
<i>Found in Translation: Reconstructing Phylogenetic Language Trees from Translations</i>	
Ella Rabinovich, Noam Ordan and Shuly Wintner	530
<i>Predicting Native Language from Gaze</i>	
Yevgeni Berzak, Chie Nakamura, Suzanne Flynn and Boris Katz	541
<i>MORSE: Semantic-ally Drive-n MORpheme SEGment-er</i>	
Tarek Sakakini, Suma Bhat and Pramod Viswanath	552

<i>Deep Pyramid Convolutional Neural Networks for Text Categorization</i> Rie Johnson and Tong Zhang	562
<i>Improved Neural Relation Detection for Knowledge Base Question Answering</i> Mo Yu, Wenpeng Yin, Kazi Saidul Hasan, Cicero dos Santos, Bing Xiang and Bowen Zhou ...	571
<i>Deep Keyphrase Generation</i> Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky and Yu Chi	582
<i>Attention-over-Attention Neural Networks for Reading Comprehension</i> Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu and Guoping Hu	593
<i>Alignment at Work: Using Language to Distinguish the Internalization and Self-Regulation Components of Cultural Fit in Organizations</i> Gabriel Doyle, Amir Goldberg, Sameer Srivastava and Michael Frank	603
<i>Representations of language in a model of visually grounded speech signal</i> Grzegorz Chrupała, Lieke Gelderloos and Afra Alishahi	613
<i>Spectral Analysis of Information Density in Dialogue Predicts Collaborative Task Performance</i> Yang Xu and David Reitter	623
<i>Affect-LM: A Neural Language Model for Customizable Affective Text Generation</i> Sayan Ghosh, Mathieu Chollet, Eugene Laksana, Louis-Philippe Morency and Stefan Scherer .	634
<i>Domain Attention with an Ensemble of Experts</i> Young-Bum Kim, Karl Stratos and Dongchan Kim	643
<i>Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders</i> Tiancheng Zhao, Ran Zhao and Maxine Eskenazi	654
<i>Hybrid Code Networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning</i> Jason D Williams, Kavosh Asadi and Geoffrey Zweig	665
<i>Generating Contrastive Referring Expressions</i> Martin Villalba, Christoph Teichmann and Alexander Koller	678
<i>Modeling Source Syntax for Neural Machine Translation</i> Junhui Li, Deyi Xiong, Zhaopeng Tu, Muhua Zhu, Min Zhang and Guodong Zhou	688
<i>Sequence-to-Dependency Neural Machine Translation</i> Shuangzhi Wu, Dongdong Zhang, Nan Yang, Mu Li and Ming Zhou	698
<i>Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning</i> Jing Ma, Wei Gao and Kam-Fai Wong	708
<i>EmoNet: Fine-Grained Emotion Detection with Gated Recurrent Neural Networks</i> Muhammad Abdul-Mageed and Lyle Ungar	718
<i>Beyond Binary Labels: Political Ideology Prediction of Twitter Users</i> Daniel Preoțiuc-Pietro, Ye Liu, Daniel Hopkins and Lyle Ungar	729

<i>Leveraging Behavioral and Social Information for Weakly Supervised Collective Classification of Political Discourse on Twitter</i>	
Kristen Johnson, Di Jin and Dan Goldwasser	741
<i>A Nested Attention Neural Hybrid Model for Grammatical Error Correction</i>	
Jianshu Ji, Qinlong Wang, Kristina Toutanova, Yongen Gong, Steven Truong and Jianfeng Gao	753
<i>TextFlow: A Text Similarity Measure based on Continuous Sequences</i>	
Yassine Mrabet, Halil Kilicoglu and Dina Demner-Fushman	763
<i>Friendships, Rivalries, and Trysts: Characterizing Relations between Ideas in Texts</i>	
Chenhao Tan, Dallas Card and Noah A. Smith	773
<i>Polish evaluation dataset for compositional distributional semantics models</i>	
Alina Wróblewska and Katarzyna Krasnowska-Kieraś	784
<i>Automatic Annotation and Evaluation of Error Types for Grammatical Error Correction</i>	
Christopher Bryant, Mariano Felice and Ted Briscoe	793
<i>Evaluation Metrics for Machine Reading Comprehension: Prerequisite Skills and Readability</i>	
Saku Sugawara, Yusuke Kido, Hikaru Yokono and Akiko Aizawa	806
<i>A Minimal Span-Based Neural Constituency Parser</i>	
Mitchell Stern, Jacob Andreas and Dan Klein	818
<i>Semantic Dependency Parsing via Book Embedding</i>	
Weiwei Sun, Junjie Cao and Xiaojun Wan	828
<i>Neural Word Segmentation with Rich Pretraining</i>	
Jie Yang, Yue Zhang and Fei Dong	839
<i>Neural Machine Translation via Binary Code Prediction</i>	
Yusuke Oda, Philip Arthur, Graham Neubig, Koichiro Yoshino and Satoshi Nakamura	850
<i>What do Neural Machine Translation Models Learn about Morphology?</i>	
Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad and James Glass	861
<i>Context-Dependent Sentiment Analysis in User-Generated Videos</i>	
Soujanya Poria, Erik Cambria, Devamanyu Hazarika, Navonil Majumder, Amir Zadeh and Louis-Philippe Morency	873
<i>A Multidimensional Lexicon for Interpersonal Stancetaking</i>	
Umashanthi Pavalanathan, Jim Fitzpatrick, Scott Kiesling and Jacob Eisenstein	884
<i>Tandem Anchoring: a Multiword Anchor Approach for Interactive Topic Modeling</i>	
Jeffrey Lund, Connor Cook, Kevin Seppi and Jordan Boyd-Graber	896
<i>Apples to Apples: Learning Semantics of Common Entities Through a Novel Comprehension Task</i>	
Omid Bakhshandeh and James Allen	906
<i>Going out on a limb: Joint Extraction of Entity Mentions and Relations without Dependency Trees</i>	
Arzoo Katiyar and Claire Cardie	917
<i>Naturalizing a Programming Language via Interactive Learning</i>	
Sida I. Wang, Samuel Ginn, Percy Liang and Christopher D. Manning	929

<i>Semantic Word Clusters Using Signed Spectral Clustering</i> Joao Sedoc, Jean Gallier, Dean Foster and Lyle Ungar	939
<i>An Interpretable Knowledge Transfer Model for Knowledge Base Completion</i> Qizhe Xie, Xuezhe Ma, Zihang Dai and Eduard Hovy	950
<i>Learning a Neural Semantic Parser from User Feedback</i> Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy and Luke Zettlemoyer .	963
<i>Joint Modeling of Content and Discourse Relations in Dialogues</i> Kechen Qin, Lu Wang and Joseph Kim	974
<i>Argument Mining with Structured SVMs and RNNs</i> Vlad Niculae, Joonsuk Park and Claire Cardie	985
<i>Neural Discourse Structure for Text Categorization</i> Yangfeng Ji and Noah A. Smith	996
<i>Adversarial Connective-exploiting Networks for Implicit Discourse Relation Classification</i> Lianhui Qin, Zhisong Zhang, Hai Zhao, Zhiting Hu and Eric Xing	1006
<i>Don't understand a measure? Learn it: Structured Prediction for Coreference Resolution optimizing its measures</i> Iryna Haponchuk and Alessandro Moschitti	1018
<i>Bayesian Modeling of Lexical Resources for Low-Resource Settings</i> Nicholas Andrews, Mark Dredze, Benjamin Van Durme and Jason Eisner	1029
<i>Semi-Supervised QA with Generative Domain-Adaptive Nets</i> Zhilin Yang, Junjie Hu, Ruslan Salakhutdinov and William Cohen	1040
<i>From Language to Programs: Bridging Reinforcement Learning and Maximum Marginal Likelihood</i> Kelvin Guu, Panupong Pasupat, Evan Liu and Percy Liang	1051
<i>Diversity driven attention model for query-based abstractive summarization</i> Preksha Nema, Mitesh M. Khapra, Anirban Laha and Balaraman Ravindran	1063
<i>Get To The Point: Summarization with Pointer-Generator Networks</i> Abigail See, Peter J. Liu and Christopher D. Manning	1073
<i>Supervised Learning of Automatic Pyramid for Optimization-Based Multi-Document Summarization</i> Maxime Peyrard and Judith Eckle-Kohler	1084
<i>Selective Encoding for Abstractive Sentence Summarization</i> Qingyu Zhou, Nan Yang, Furu Wei and Ming Zhou	1095
<i>PositionRank: An Unsupervised Approach to Keyphrase Extraction from Scholarly Documents</i> Corina Florescu and Cornelia Caragea	1105
<i>Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses</i> Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio and Joelle Pineau	1116
<i>A Transition-Based Directed Acyclic Graph Parser for UCCA</i> Daniel Hershcovich, Omri Abend and Ari Rappoport	1127

<i>Abstract Syntax Networks for Code Generation and Semantic Parsing</i> Maxim Rabinovich, Mitchell Stern and Dan Klein	1139
<i>Visualizing and Understanding Neural Machine Translation</i> Yanzhuo Ding, Yang Liu, Huanbo Luan and Maosong Sun	1150
<i>Detecting annotation noise in automatically labelled data</i> Ines Rehbein and Josef Ruppenhofer	1160
<i>Abstractive Document Summarization with a Graph-Based Attentional Neural Model</i> Jiwei Tan, Xiaojun Wan and Jianguo Xiao	1171
<i>Probabilistic Typology: Deep Generative Models of Vowel Inventories</i> Ryan Cotterell and Jason Eisner	1182
<i>Adversarial Multi-Criteria Learning for Chinese Word Segmentation</i> Xinchi Chen, Zhan Shi, Xipeng Qiu and Xuanjing Huang	1193
<i>Neural Joint Model for Transition-based Chinese Syntactic Analysis</i> Shuhe Kurita, Daisuke Kawahara and Sadao Kurohashi	1204
<i>Robust Incremental Neural Semantic Graph Parsing</i> Jan Buys and Phil Blunsom	1215
<i>Joint Extraction of Entities and Relations Based on a Novel Tagging Scheme</i> Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou and Bo Xu	1227
<i>A Local Detection Approach for Named Entity Recognition and Mention Detection</i> Mingbin Xu, Hui Jiang and Sedtawut Watcharawittayakul	1237
<i>Vancouver Welcomes You! Minimalist Location Metonymy Resolution</i> Milan Gritta, Mohammad Taher Pilehvar, Nut Limsopatham and Nigel Collier	1248
<i>Unifying Text, Metadata, and User Network Representations with a Neural Network for Geolocation Prediction</i> Yasuhide Miura, Motoki Taniguchi, Tomoki Taniguchi and Tomoko Ohkuma	1260
<i>Multi-Task Video Captioning with Video and Entailment Generation</i> Ramakanth Pasunuru and Mohit Bansal	1273
<i>Enriching Complex Networks with Word Embeddings for Detecting Mild Cognitive Impairment from Speech Transcripts</i> Leandro Santos, Edilson Anselmo Corrêa Júnior, Osvaldo Oliveira Jr, Diego Amancio, Letícia Mansur and Sandra Aluísio	1284
<i>Adversarial Adaptation of Synthetic or Stale Data</i> Young-Bum Kim, Karl Stratos and Dongchan Kim	1297
<i>Chat Detection in an Intelligent Assistant: Combining Task-oriented and Non-task-oriented Spoken Dialogue Systems</i> Satoshi Akasaki and Nobuhiro Kaji	1308
<i>A Neural Local Coherence Model</i> Dat Tien Nguyen and Shafiq Joty	1320

<i>Data-Driven Broad-Coverage Grammars for Opinionated Natural Language Generation (ONLG)</i> Tomer Cagan, Stefan L. Frank and Reut Tsarfaty	1331
<i>Learning to Ask: Neural Question Generation for Reading Comprehension</i> Xinya Du, Junru Shao and Claire Cardie	1342
<i>Joint Optimization of User-desired Content in Multi-document Summaries by Learning from User Feedback</i> Avinesh PVS and Christian M. Meyer	1353
<i>Flexible and Creative Chinese Poetry Generation Using Neural Memory</i> Jiyuan Zhang, Yang Feng, Dong Wang, Yang Wang, Andrew Abel, Shiyue Zhang and Andi Zhang	1364
<i>Learning to Generate Market Comments from Stock Prices</i> Soichiro Murakami, Akihiko Watanabe, Akira Miyazawa, Keiichi Goshima, Toshihiko Yanase, Hiroya Takamura and Yusuke Miyao	1374
<i>Can Syntax Help? Improving an LSTM-based Sentence Compression Model for New Domains</i> Liangguo Wang, Jing Jiang, Hai Leong Chieu, Chen Hui Ong, Dandan Song and Lejian Liao	1385
<i>Transductive Non-linear Learning for Chinese Hypernym Prediction</i> Chengyu Wang, Junchi Yan, Aoying Zhou and Xiaofeng He	1394
<i>A Constituent-Centric Neural Architecture for Reading Comprehension</i> Pengtao Xie and Eric Xing	1405
<i>Cross-lingual Distillation for Text Classification</i> Ruochen Xu and Yiming Yang	1415
<i>Understanding and Predicting Empathic Behavior in Counseling Therapy</i> Verónica Pérez-Rosas, Rada Mihalcea, Kenneth Resnicow, Satinder Singh and Lawrence An	1426
<i>Leveraging Knowledge Bases in LSTMs for Improving Machine Reading</i> Bishan Yang and Tom Mitchell	1436
<i>Prerequisite Relation Learning for Concepts in MOOCs</i> Liangming Pan, Chengjiang Li, Juanzi Li and Jie Tang	1447
<i>Unsupervised Text Segmentation Based on Native Language Characteristics</i> Shervin Malmasi, Mark Dras, Mark Johnson, Lan Du and Magdalena Wolska	1457
<i>Weakly Supervised Cross-Lingual Named Entity Recognition via Effective Annotation and Representation Projection</i> Jian Ni, Georgiana Dinu and Radu Florian	1470
<i>Context Sensitive Lemmatization Using Two Successive Bidirectional Gated Recurrent Networks</i> Abhisek Chakrabarty, Onkar Arun Pandit and Utpal Garain	1481
<i>Learning to Create and Reuse Words in Open-Vocabulary Neural Language Modeling</i> Kazuya Kawakami, Chris Dyer and Phil Blunsom	1492
<i>Bandit Structured Prediction for Neural Sequence-to-Sequence Learning</i> Julia Kreutzer, Artem Sokolov and Stefan Riezler	1503

<i>Prior Knowledge Integration for Neural Machine Translation using Posterior Regularization</i>	
Jiacheng Zhang, Yang Liu, Huanbo Luan, Jingfang Xu and Maosong Sun	1514
<i>Incorporating Word Reordering Knowledge into Attention-based Neural Machine Translation</i>	
Jinchao Zhang, Mingxuan Wang, Qun Liu and Jie Zhou	1524
<i>Lexically Constrained Decoding for Sequence Generation Using Grid Beam Search</i>	
Chris Hokamp and Qun Liu	1535
<i>Combating Human Trafficking with Multimodal Deep Models</i>	
Edmund Tong, Amir Zadeh, Cara Jones and Louis-Philippe Morency	1547
<i>MalwareTextDB: A Database for Annotated Malware Articles</i>	
Swee Kiat Lim, Aldrian Obaja Muis, Wei Lu and Chen Hui Ong	1557
<i>A Corpus of Annotated Revisions for Studying Argumentative Writing</i>	
Fan Zhang, Homa B. Hashemi, Rebecca Hwa and Diane Litman	1568
<i>Automatic Induction of Synsets from a Graph of Synonyms</i>	
Dmitry Ustalov, Alexander Panchenko and Chris Biemann	1579
<i>Neural Modeling of Multi-Predicate Interactions for Japanese Predicate Argument Structure Analysis</i>	
Hiroki Ouchi, Hiroyuki Shindo and Yuji Matsumoto	1591
<i>TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension</i>	
Mandar Joshi, Eunsol Choi, Daniel Weld and Luke Zettlemoyer	1601
<i>Learning Semantic Correspondences in Technical Documentation</i>	
Kyle Richardson and Jonas Kuhn	1612
<i>Bridge Text and Knowledge by Learning Multi-Prototype Entity Mention Embedding</i>	
Yixin Cao, Lifu Huang, Heng Ji, Xu Chen and Juanzi Li	1623
<i>Interactive Learning of Grounded Verb Semantics towards Human-Robot Communication</i>	
Lanbo She and Joyce Chai	1634
<i>Multimodal Word Distributions</i>	
Ben Athiwaratkun and Andrew Wilson	1645
<i>Enhanced LSTM for Natural Language Inference</i>	
Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang and Diana Inkpen	1657
<i>Linguistic analysis of differences in portrayal of movie characters</i>	
Anil Ramakrishna, Victor R. Martínez, Nikolaos Malandrakis, Karan Singla and Shrikanth Narayanan	1669
<i>Linguistically Regularized LSTM for Sentiment Classification</i>	
Qiao Qian, Minlie Huang, Jinhao Lei and Xiaoyan Zhu	1679
<i>Sarcasm SIGN: Interpreting Sarcasm with Sentiment Based Monolingual Machine Translation</i>	
Lotem Peled and Roi Reichart	1690
<i>Active Sentiment Domain Adaptation</i>	
Fangzhao Wu, Yongfeng Huang and Jun Yan	1701

<i>Volatility Prediction using Financial Disclosures Sentiments with Word Embedding-based IR Models</i> Navid Rekabsaz, Mihai Lupu, Artem Baklanov, Alexander Dür, Linda Andersson and Allan Hanbury	1712
<i>CANE: Context-Aware Network Embedding for Relation Modeling</i> Cunchao Tu, Han Liu, Zhiyuan Liu and Maosong Sun	1722
<i>Universal Dependencies Parsing for Colloquial Singaporean English</i> Hongmin Wang, Yue Zhang, GuangYong Leonard Chan, Jie Yang and Hai Leong Chieu	1732
<i>Generic Axiomatization of Families of Noncrossing Graphs in Dependency Parsing</i> Anssi Yli-Jyrä and Carlos Gómez-Rodríguez	1745
<i>Semi-supervised sequence tagging with bidirectional language models</i> Matthew Peters, Waleed Ammar, Chandra Bhagavatula and Russell Power	1756
<i>Learning Symmetric Collaborative Dialogue Agents with Dynamic Knowledge Graph Embeddings</i> He He, Anusha Balakrishnan, Mihail Eric and Percy Liang	1766
<i>Neural Belief Tracker: Data-Driven Dialogue State Tracking</i> Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson and Steve Young	1777
<i>Exploiting Argument Information to Improve Event Detection via Supervised Attention Mechanisms</i> Shulin Liu, Yubo Chen, Kang Liu and Jun Zhao	1789
<i>Topical Coherence in LDA-based Models through Induced Segmentation</i> Hesam Amoualian, Wei Lu, Eric Gaussier, Georgios Balikas, Massih R Amini and Marianne Clausel	1799
<i>Jointly Extracting Relations with Class Ties via Effective Deep Ranking</i> Hai Ye, Wenhan Chao, Zhunchen Luo and Zhoujun Li	1810
<i>Search-based Neural Structured Learning for Sequential Question Answering</i> Mohit Iyyer, Wen-tau Yih and Ming-Wei Chang	1821
<i>Gated-Attention Readers for Text Comprehension</i> Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen and Ruslan Salakhutdinov	1832
<i>Determining Gains Acquired from Word Embedding Quantitatively Using Discrete Distribution Clustering</i> Jianbo Ye, Yanran Li, Zhaohui Wu, James Z. Wang, Wenjie Li and Jia Li	1847
<i>Towards a Seamless Integration of Word Senses into Downstream NLP Applications</i> Mohammad Taher Pilehvar, Jose Camacho-Collados, Roberto Navigli and Nigel Collier	1857
<i>Reading Wikipedia to Answer Open-Domain Questions</i> Danqi Chen, Adam Fisch, Jason Weston and Antoine Bordes	1870
<i>Learning to Skim Text</i> Adams Wei Yu, Hongrae Lee and Quoc Le	1880
<i>An Algebra for Feature Extraction</i> Vivek Srikumar	1891

<i>Chunk-based Decoder for Neural Machine Translation</i>	
Shonosuke Ishiwatari, Jingtao Yao, Shujie Liu, Mu Li, Ming Zhou, Naoki Yoshinaga, Masaru Kitsuregawa and Weijia Jia	1901
<i>Doubly-Attentive Decoder for Multi-modal Neural Machine Translation</i>	
Iacer Calixto, Qun Liu and Nick Campbell	1913
<i>A Teacher-Student Framework for Zero-Resource Neural Machine Translation</i>	
Yun Chen, Yang Liu, Yong Cheng and Victor O.K. Li	1925
<i>Improved Neural Machine Translation with a Syntax-Aware Encoder and Decoder</i>	
Huadong Chen, Shujian Huang, David Chiang and Jiajun Chen	1936
<i>Cross-lingual Name Tagging and Linking for 282 Languages</i>	
Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight and Heng Ji	1946
<i>Adversarial Training for Unsupervised Bilingual Lexicon Induction</i>	
Meng Zhang, Yang Liu, Huanbo Luan and Maosong Sun	1959
<i>Estimating Code-Switching on Twitter with a Novel Generalized Word-Level Language Detection Technique</i>	
Shruti Rijhwani, Royal Sequiera, Monojit Choudhury, Kalika Bali and Chandra Shekhar Maddila	1971
<i>Using Global Constraints and Reranking to Improve Cognates Detection</i>	
Michael Bloodgood and Benjamin Strauss	1983
<i>One-Shot Neural Cross-Lingual Transfer for Paradigm Completion</i>	
Katharina Kann, Ryan Cotterell and Hinrich Schütze	1993
<i>Morphological Inflection Generation with Hard Monotonic Attention</i>	
Roei Aharoni and Yoav Goldberg	2004
<i>From Characters to Words to in Between: Do We Capture Morphology?</i>	
Clara Vania and Adam Lopez	2016
<i>Riemannian Optimization for Skip-Gram Negative Sampling</i>	
Alexander Fonarev, Oleksii Grinchuk, Gleb Gusev, Pavel Serdyukov and Ivan Oseledets	2028
<i>Deep Multitask Learning for Semantic Dependency Parsing</i>	
Hao Peng, Sam Thomson and Noah A. Smith	2037
<i>Improved Word Representation Learning with Sememes</i>	
Yilin Niu, Ruobing Xie, Zhiyuan Liu and Maosong Sun	2049
<i>Learning Character-level Compositionality with Visual Features</i>	
Frederick Liu, Han Lu, Chieh Lo and Graham Neubig	2059
<i>A Progressive Learning Approach to Chinese SRL Using Heterogeneous Data</i>	
Qiaolin Xia, Lei Sha, Baobao Chang and Zhifang Sui	2069
<i>Revisiting Recurrent Networks for Paraphrastic Sentence Embeddings</i>	
John Wieting and Kevin Gimpel	2078
<i>Ontology-Aware Token Embeddings for Prepositional Phrase Attachment</i>	
Pradeep Dasigi, Waleed Ammar, Chris Dyer and Eduard Hovy	2089

<i>Identifying 1950s American Jazz Musicians: Fine-Grained IsA Extraction via Modifier Composition</i>	
Ellie Pavlick and Marius Pasca	2099
<i>Parsing to 1-Endpoint-Crossing, Pagenumber-2 Graphs</i>	
Junjie Cao, Sheng Huang, Weiwei Sun and Xiaojun Wan	2110
<i>Semi-supervised Multitask Learning for Sequence Labeling</i>	
Marek Rei	2121
<i>Semantic Parsing of Pre-university Math Problems</i>	
Takuya Matsuzaki, Takumi Ito, Hidenao Iwane, Hirokazu Anai and Noriko H. Arai	2131