

Accepted papers - ACL 2022

[Main Conference](#)

[Long Papers](#)

[Short Papers](#)

[System Demonstration Papers](#)

[FINDINGS OF ACL 2022](#)

[Long Papers](#)

[Short Papers](#)

Main Conference

Long Papers

Title	Authors
A Closer Look at How Fine-tuning Changes BERT	Yichu Zhou, Vivek Srikumar
A Comparative Study of Faithfulness Metrics for Model Interpretability Methods	Chun Sik Chan, Huanqi Kong, LIANG GUANQING
A Comparison of Strategies for Source-Free Domain Adaptation	Xin Su, Yiyun Zhao, Steven Bethard
A Contrastive Framework for Learning Sentence Representations from Pairwise and Triple-wise Perspective in Angular Space	Yuhao Zhang, Hongji Zhu, Yongliang Wang, Nan Xu, XIAOBO LI, Binqiang Zhao
A Functionalist Account of Vowel System Typology	Ryan D Cotterell, Jason Eisner
A Good Prompt Is Worth Millions of Parameters: Low-resource Prompt-based Learning for Vision-Language Models	Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen, Xiang Ren
A Meta-framework for Spatiotemporal Quantity Extraction from Text	Qiang Ning, Ben Zhou, Hao Wu, Haoruo Peng, Chuchu Fan, Matt Gardner
A Model-agnostic Data Manipulation Method for Persona-based Dialogue Generation	Yu Cao, Wei Bi, Meng Fang, Shuming Shi, Dacheng Tao
A Multi-Document Coverage Reward for RELAXed Multi-Document Summarization	Jacob Parnell, Inigo Jauregi Unanue, Massimo Piccardi
A Neural Network Architecture for Program Understanding Inspired by Human Behaviors	Renyu Zhu, Lei Yuan, Xiang Li, Ming Gao, Wenyan Cai
A Rationale-Centric Framework for Human-in-the-loop Machine Learning	JINGHUI LU, Linyi Yang, Brian Mac Namee, Yue Zhang
A Statutory Article Retrieval Dataset in French	Antoine Louis, Gerasimos Spanakis
A Taxonomy of Empathetic Questions in Social Dialogs	Ekaterina Svikhnushina, Iuliana Voinea, Anuradha Welivita, Pearl Pu

A Token-level Reference-free Hallucination Detection Benchmark for Free-form Text Generation	Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, Bill Dolan
A Variational Hierarchical Model for Neural Cross-Lingual Summarization	Yunlong Liang, Fandong Meng, Chulun Zhou, Jinan Xu, Yufeng Chen, Jinsong Su, Jie Zhou
A Well-Composed Text is Half Done! Composition Sampling for Diverse Conditional Generation	Shashi Narayan, Gonçalo Simões, Yao Zhao, Joshua Maynez, Dipanjan Das, Michael Collins, Mirella Lapata
ABC: Attention with Bounded-memory Control	Hao Peng, Jungo Kasai, Nikolaos Pappas, Dani Yogatama, Zhaofeng Wu, Lingpeng Kong, Roy Schwartz, Noah Smith
Accelerating Code Search with Deep Hashing and Code Classification	Wenchao Gu, Yanlin Wang, Lun Du, Hongyu Zhang, Shi Han, Dongmei Zhang, Michael Lyu
Accurate Online Posterior Alignments for Principled Lexically-Constrained Decoding	Soumya Chatterjee, Sunita Sarawagi, Preethi Jyothi
Achieving Conversational Goals with Unsupervised Post-hoc Knowledge Injection	Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, Julian McAuley
Achieving Reliable Human Assessment of Open-Domain Dialogue Systems	Tianbo Ji, Yvette Graham, Gareth J. F. Jones, Chenyang Lyu, Qun Liu
Active Evaluation: Efficient NLG Evaluation with Few Pairwise Comparisons	Akash Kumar Mohankumar, Mitesh M Khapra
AdaLoGN: Adaptive Logic Graph Network for Reasoning-Based Machine Reading Comprehension	Xiao Li, Gong Cheng, Ziheng Chen, Yawei Sun, Yuzhong Qu
AdapLeR: Speeding up Inference by Adaptive Length Reduction	Ali Modarressi, Hosein Mohebbi, Mohammad Taher Pilehvar
Adapting Coreference Resolution Models through Active Learning	Michelle Yuan, Patrick Xia, Chandler May, Benjamin Van Durme, Jordan Lee Boyd-Graber
Adaptive Testing and Debugging of NLP Models	Marco Tulio Ribeiro, Scott M Lundberg
Adversarial Authorship Attribution for Deobfuscation	Wanyue Zhai, Jonathan Rusert, Zubair Shafiq, Padmini Srinivasan
Adversarial Soft Prompt Tuning for Cross-Domain Sentiment Analysis	Hui Wu, Xiaodong Shi
AlephBERT: Language Model Pre-training and Evaluation from Sub-Word to Sentence Level	Amit Seker, Elron Bandel, Dan Bareket, Idan Brusilovsky, Refael Shaked Greenfeld, Reut Tsarfaty
Alignment-Augmented Consistent Translation for Multilingual Open Information Extraction	Keshav Kolluru, Muqeeth Mohammed, Shubham Mittal, Soumen Chakrabarti, Mausam .
Alternative Input Signals Ease Transfer in Multilingual Machine Translation	Simeng Sun, Angela Fan, James Cross, Vishrav Chaudhary, Chau Tran, Philipp Koehn, Francisco Guzmán
AmericasNLI: Evaluating Zero-shot Natural Language Understanding of Pretrained Multilingual Models in Truly Low-resource Languages	Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Vladimir Meza Ruiz, Gustavo A. Giménez-Lugo,

	Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Thang Vu, Katharina Kann
An Effective and Efficient Entity Alignment Decoding Algorithm via Third-Order Tensor Isomorphism	Xin Mao, Meirong Ma, Hao Yuan, Jianchao Zhu, ZongYu Wang, Rui Xie, Wei Wu, Man Lan
An Empirical Study of Memorization in NLP	Xiaosen Zheng, Jing Jiang
An Empirical Study on Explanations in Out-of-Domain Settings	George Chrysostomou, Nikolaos Aletras
An Empirical Survey of the Effectiveness of Debiasing Techniques for Pre-trained Language Models	Nicholas Meade, Elinor Poole-Dayan, Siva Reddy
An Imitation Learning Curriculum for Text Editing with Non-Autoregressive Models	Sweta Agrawal, Marine Carpuat
An Information-theoretic Approach to Prompt Engineering Without Ground Truth Labels	Taylor Sorensen, Joshua Robinson, Christopher Michael Rytting, Alexander Glenn Shaw, Kyle Jeffrey Rogers, Alexia Pauline Delorey, Mahmoud Khalil, Nancy Fulda, David Wingate
An Interpretable Neuro-Symbolic Reasoning Framework for Task-Oriented Dialogue Generation	Shiquan Yang, Rui Zhang, Sarah Monazam Erfani, Jey Han Lau
An Investigation of the (In)effectiveness of Counterfactually Augmented Data	Nitish Joshi, He He
An Unsupervised Multiple-Task and Multiple-Teacher Model for Cross-lingual Named Entity Recognition	Zhuoran Li, Chunming Hu, Xiaohui Guo, Junfan Chen, Wenyi Qin, Richong Zhang
Analyzing Generalization of Vision and Language Navigation to Unseen Outdoor Areas	Raphael Schumann, Stefan Riezler
Answer-level Calibration for Free-form Multiple Choice Question Answering	Sawan Kumar
Answering Open-Domain Multi-Answer Questions via a Recall-then-Verify Framework	Zhihong Shao, Minlie Huang
AraT5: Text-to-Text Transformers for Arabic Language Generation	El Moatez Billah Nagoudi, AbdelRahim A. Elmadany, Muhammad Abdul-Mageed
Are Prompt-based Models Clueless?	Pride Kavumba, Ryo Takahashi, Yusuke Oda
ASPECTNEWS: Aspect-Oriented Summarization of News Documents	Ojas Ahuja, Jiacheng Xu, Akshay Kumar Gupta, Kevin Horecka, Greg Durrett
Attention Temperature Matters in Abstractive Summarization Distillation	Shengqiang Zhang, Xingxing Zhang, Hangbo Bao, Furu Wei
Auto-Debias: Debiasing Masked Language Models with Automated Biased Prompts	Yue Guo, Yi Yang, Ahmed Abbasi
Automated Crossword Solving	Eric Wallace, Nicholas Tomlin, Albert Xu, Kevin Yang, Eshaan Pathak, Matthew L. Ginsberg, Dan Klein
Automatic Error Analysis for Document-level Information Extraction	Aliva Das, Xinya Du, Barry Wang, Kejian Shi, Jiayuan Gu, Thomas Porter, Claire Cardie

Automatic Identification and Classification of Bragging in Social Media	Mali Jin, Daniel Preotiuc-Pietro, A. Seza Dođruöz, Nikolaos Aletras
Bag-of-Words vs. Graph vs. Sequence in Text Classification: Questioning the Necessity of Text-Graphs and the Surprising Strength of a Wide MLP	Lukas Galke, Ansgar Scherp
BenchIE: A Framework for Multi-Faceted Fact-Based Open Information Extraction Evaluation	Kiril Gashteovski, Mingying Yu, Bhushan Kotnis, Carolin Lawrence, Mathias Niepert, Goran Glavaš
BERT Learns to Teach: Knowledge Distillation with Meta Learning	Wangchunshu Zhou, Canwen Xu, Julian McAuley
bert2BERT: Towards Reusable Pretrained Language Models	Cheng Chen, Yichun Yin, Lifeng Shang, Xin Jiang, Yujia Qin, Fengyu Wang, Zhi Wang, Xiao Chen, Zhiyuan Liu, Qun Liu
Better Language Model with Hypernym Class Prediction	He Bai, Tong Wang, Alessandro Sordoni, Peng Shi
Beyond Goldfish Memory: Long-Term Open-Domain Conversation	Jing Xu, Arthur Szlam, Jason E Weston
Beyond the Granularity: Multi-Perspective Dialogue Collaborative Selection for Dialogue State Tracking	Jinyu Guo, Kai Shuang, Jijie Li, Zihan Wang, Yixuan Liu
Bias Mitigation in Machine Translation Quality Estimation	Hanna Behnke, Marina Fomicheva, Lucia Specia
Bilingual alignment transfers to multilingual alignment for unsupervised parallel text mining	Chih-chan Tien, Shane Steinert-Threlkeld
BITIIMT: A Bilingual Text-infilling Method for Interactive Machine Translation	Yanling Xiao, Lemao Liu, Guoping Huang, Qu Cui, Shujian Huang, Shuming Shi, Jiajun Chen
Bottom-Up Constituency Parsing and Nested Named Entity Recognition with Pointer Networks	Songlin Yang, Kewei Tu
Boundary Smoothing for Named Entity Recognition	Enwei Zhu, Jinpeng Li
Bridging the Data Gap between Training and Inference for Unsupervised Neural Machine Translation	Zhiwei He, Xing Wang, Rui Wang, Shuming Shi, Zhaopeng Tu
Bridging the Generalization Gap in Text-to-SQL Parsing with Schema Expansion	Chen Zhao, Yu Su, Adam Pauls, Emmanouil Antonios Platanios
BRIO: Bringing Order to Abstractive Summarization	Yixin Liu, Pengfei Liu, Dragomir Radev, Graham Neubig
CAKE: A Scalable Commonsense-Aware Framework For Multi-View Knowledge Graph Completion	Guanglin Niu, Bo Li, Yongfei Zhang, Shiliang Pu
CaMEL: Case Marker Extraction without Labels	Leonie Weissweiler, Valentin Hofmann, Masoud Jalili Sabet, Hinrich Schuetze
CAMERO: Consistency Regularized Ensemble of Perturbed Language Models with Weight Sharing	Chen Liang, Pengcheng He, Yelong Shen, Weizhu Chen, Tuo Zhao
Can Explanations Be Useful for Calibrating Black Box Models?	Xi Ye, Greg Durrett
Can Pre-trained Language Models Interpret Similes as Smart as Human?	Qianyu He, Sijie Cheng, Zhixu Li, Rui Xie, Yanghua Xiao

Can Prompt Probe Pretrained Language Models? Understanding the Invisible Risks from a Causal View	Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu, Le Sun
Can Synthetic Translations Improve Bitext Quality?	Eleftheria Briakou, Marine Carpuat
Can Transformer be Too Compositional? Analysing Idiom Processing in Neural Machine Translation	Verna Dankers, Christopher G. Lucas, Ivan Titov
Can Unsupervised Knowledge Transfer from Social Discussions Help Argument Mining?	Subhabrata Dutta, Jeevesh Juneja, Dipankar Das, Tanmoy Chakraborty
CARETS: A Consistency And Robustness Evaluative Test Suite for VQA	Carlos E Jimenez, Olga Russakovsky, Karthik R Narasimhan
CASPI Causal-aware Safe Policy Improvement for Task-oriented Dialogue	Govardana Sachithanandam Ramachandran, Kazuma Hashimoto, Caiming Xiong
Causal Probing for Grammatical Number: From Encoding to Usage	Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau, Ryan D Cotterell
CBLUE: A Chinese Biomedical Language Understanding Evaluation Benchmark	Ningyu Zhang, Mosha Chen, Zhen Bi, Xiaozhuan Liang, Lei Li, Xin Shang, Kangping Yin, Chuanqi Tan, Jian Xu, Fei Huang, Luo Si, Yuan Ni, Guotong Xie, Zhifang Sui, Baobao Chang, Hui Zong, Zheng Yuan, Linfeng Li, Jun Yan, Hongying ZAN, Kunli Zhang, Buzhou Tang, Qingcai Chen
Challenges and Strategies in Cross-Cultural NLP	Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, Anders Søgaard
Characterizing Idioms: Conventionality and Contingency	Michaela Socolof, Jackie CK Cheung, Michael Wagner, Timothy J. O'Donnell
Chart-to-Text: A Large-Scale Benchmark for Chart Summarization	Shankar Kantharaj, Rixie Tiffany Ko Leong, Xiang Lin, Ahmed Masry, Megh Thakkar, Enamul Hoque, Shafiq Joty
ChatMatch: Evaluating Chatbots by Autonomous Chat Tournaments	Ruolan Yang, Zitong Li, Haifeng Tang, Kenny Q. Zhu
CICERO: A Dataset for Contextualized Commonsense Inference in Dialogues	Deepanway Ghosal, Siqi Shen, Navonil Majumder, Rada Mihalcea, Soujanya Poria
CipherDAug: Ciphertext based Data Augmentation for Neural Machine Translation	Nishant Kambhatla, Logan Born, Anoop Sarkar
ClarET: Pre-training a Correlation-Aware Context-To-Event Transformer for Event-Centric Generation and Classification	Yucheng Zhou, Tao Shen, Xiubo Geng, Guodong Long, Daxin Jiang
Clickbait Spoiling via Question Answering and Passage Retrieval	Matthias Hagen, Maik Fröbe, Artur Jurk, Martin Potthast
CLIP Models are Few-Shot Learners: Empirical Studies on VQA and Visual Entailment	Haoyu Song, Li Dong, Weinan Zhang, Ting Liu, Furu Wei

CLUES: A Benchmark for Learning Classifiers using Natural Language Explanations	Rakesh R Menon, Sayan Ghosh, Shashank Srivastava
Cluster & Tune: Boost Cold Start Performance in Text Classification	Eyal Shnarch, Ariel Gera, Alon Halfon, Lena Dankin, Leshem Choshen, Ranit Aharonov, Noam Slonim
ClusterFormer: Neural Clustering Attention for Efficient and Effective Transformer	Ningning Wang, Guobing Gan, Peng Zhang, Shuai Zhang, Junqiu Wei, Qun Liu, Xin Jiang
CogTaskonomy: Cognitively Inspired Task Taxonomy Is Beneficial to Transfer Learning in NLP	yifei luo, Minghui Xu, Deyi Xiong
Coherence boosting: When your pretrained language model is not paying enough attention	Nikolay Malkin, Zhen Wang, Nebojsa Jojic
Compact Token Representations with Contextual Quantization for Efficient Document Re-ranking	Yingrui Yang, Yifan Qiao, Tao Yang
Composable Sparse Fine-Tuning for Cross-Lingual Transfer	Alan Ansell, Edoardo Ponti, Anna Korhonen, Ivan Vulić
Compositional Generalization in Dependency Parsing	Emily Goodwin, Siva Reddy, Timothy J. O'Donnell, Dzmitry Bahdanau
Compression of Generative Pre-trained Language Models via Quantization	Chaofan Tao, Lu Hou, Wei Zhang, Lifeng Shang, Xin Jiang, Qun Liu, Ping Luo, Ngai Wong
Computational Historical Linguistics and Language Diversity in South Asia	Aryaman Arora, Adam Farris, Samopriya Basu, Suresh Kolichala
Conditional Bilingual Mutual Information Based Adaptive Training for Neural Machine Translation	Songming Zhang, Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu, Jian Liu, Jie Zhou
ConditionalQA: A Complex Reading Comprehension Dataset with Conditional Answers	Haitian Sun, William W. Cohen, Ruslan Salakhutdinov
Confidence Based Bidirectional Global Context Aware Training Framework for Neural Machine Translation	Chulun Zhou, Fandong Meng, Jie Zhou, Min Zhang, Hongji Wang, Jinsong Su
Constrained Multi-Task Learning for Bridging Resolution	Hideo Kobayashi, Yufang Hou, Vincent Ng
CONTaiNER: Few-Shot Named Entity Recognition via Contrastive Learning	Sarkar Snigdha Sarathi Das, Arzoo Katiyar, Rebecca J. Passonneau, Rui Zhang
Context Matters: A Pragmatic Study of PLMs' Negation Understanding	Reto Gubelmann, Siegfried Handschuh
Contextual Fine-to-Coarse Distillation for Coarse-grained Response Selection in Open-Domain Conversations	Wei Chen, Yeyun Gong, Can Xu, Huang Hu, Bolun Yao, zhongyu wei, Zhihao Fan, Xiaowu Hu, Bartuer Zhou, Biao Cheng, Daxin Jiang, Nan Duan
Contextual Representation Learning beyond Masked Language Modeling	Zhiyi Fu, Wangchunshu Zhou, Jingjing Xu, Hao Zhou, Lei Li
ConTinTin: Continual Learning from Task Instructions	Wenpeng Yin, Jia Li, Caiming Xiong
Continual Few-shot Relation Learning via Embedding Space Regularization and Data Augmentation	Chengwei Qin, Shafiq Joty

Continual Pre-training of Language Models for Math Problem Understanding with Syntax-Aware Memory Network	Zheng Gong, Kun Zhou, Xin Zhao, Jing Sha, Shijin Wang, Ji-Rong Wen
Continual Prompt Tuning for Dialog State Tracking	Qi Zhu, Bing Li, Fei Mi, Xiaoyan Zhu, Minlie Huang
Continual Sequence Generation with Adaptive Compositional Modules	Yanzhe Zhang, Xuezhi Wang, Diyi Yang
Contrastive Visual Semantic Pretraining Magnifies the Semantics of Natural Language Representations	Robert Wolfe, Aylin Caliskan
Controllable Dictionary Example Generation: Generating Example Sentences for Specific Targeted Audiences	Xingwei He, Siu Ming Yiu
CQG: A Simple and Effective Controlled Generation Framework for Multi-hop Question Generation	Zichu Fei, Qi Zhang, Tao Gui, Di Liang, Sirui Wang, Wei Wu, Xuanjing Huang
Cree Corpus: A Collection of nêhiyawêwin Resources	Daniela Teodorescu, Josie Matalaski, Delaney Alexa Lothian, Denilson Barbosa, Carrie Demmans Epp
Cross-Lingual Ability of Multilingual Masked Language Models: A Study of Language Structure	Yuan Chai, Yaobo Liang, Nan Duan
Cross-Lingual Contrastive Learning for Fine-Grained Entity Typing for Low-Resource Languages	Xu Han, Yuqi Luo, Weize Chen, Zhiyuan Liu, Maosong Sun, Zhou Botong, Hao Fei, Suncong Zheng
Cross-Lingual Phrase Retrieval	Heqi Zheng, Xiao Zhang, Zewen Chi, Heyan Huang, Yan Tan, Tian Lan, Wei Wei, Xian-Ling Mao
Cross-Modal Discrete Representation Learning	Alexander H. Liu, SouYoung Jin, Cheng-I Lai, Andrew Rouditchenko, Aude Oliva, James R. Glass
Cross-Task Generalization via Natural Language Crowdsourcing Instructions	Swaroop Mishra, Daniel Khashabi, Chitta Baral, Hannaneh Hajishirzi
Cross-Utterance Conditioned VAE for Non-Autoregressive Text-to-Speech	Yang Li, Cheng Yu, Guangzhi Sun, Hua Jiang, Fanglei Sun, Weiqin Zu, Ying Wen, Yang Yang, Jun Wang
CTRL Eval: An Unsupervised Reference-Free Metric for Evaluating Controlled Text Generation	Pei Ke, Hao Zhou, Yankai Lin, Peng Li, Jie Zhou, Xiaoyan Zhu, Minlie Huang
Dataset Geography: Mapping Language Data to Language Users	Fahim Faisal, Yinkai Wang, Antonios Anastasopoulos
De-Bias for Generative Extraction in Unified NER Task	Shuai Zhang, Yongliang Shen, Zeqi Tan, Yiquan Wu, Weiming Lu
DEAM: Dialogue Coherence Evaluation using AMR-based Semantic Manipulations	Sarik Ghazarian, Nuan Wen, Aram Galstyan, Nanyun Peng
Debiased Contrastive Learning of Unsupervised Sentence Representations	Kun Zhou, Beichen Zhang, Xin Zhao, Ji-Rong Wen
Decoding Part-of-Speech from Human EEG Signals	Alex Graeme Murphy, Bernd Bohnet, Ryan

	McDonald, Uta Noppeneay
Deduplicating Training Data Makes Language Models Better	Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, Nicholas Carlini
Deep Inductive Logic Reasoning for Multi-Hop Reading Comprehension	Wenya Wang, Sinno Pan
DEEP: DEnoising Entity Pre-training for Neural Machine Translation	Junjie Hu, Hiroaki Hayashi, Kyunghyun Cho, Graham Neubig
Dependency Parsing as MRC-based Span-Span Prediction	Leilei Gan, Yuxian Meng, Kun Kuang, Xiaofei Sun, Chun Fan, Fei Wu, Jiwei Li
Dependency-based Mixture Language Models	Zhixian Yang, Xiaojun Wan
Detecting Unassimilated Borrowings in Spanish: An Annotated Corpus and Approaches to Modeling	Elena Álvarez-Mellado, Constantine Lignos
DialFact: A Benchmark for Fact-Checking in Dialogue	Prakhar Gupta, Chien-Sheng Wu, Wenhao Liu, Caiming Xiong
DialogVED: A Pre-trained Latent Variable Encoder-Decoder Model for Dialog Response Generation	Wei Chen, Yeyun Gong, Song Wang, Bolun Yao, Weizhen Qi, zhongyu wei, Xiaowu Hu, Bartuer Zhou, Yi Mao, Weizhu Chen, Biao Cheng, Nan Duan
DiBiMT: A Novel Benchmark for Measuring Word Sense Disambiguation Biases in Machine Translation	Niccolò Campolungo, Federico Martelli, Francesco Saina, Roberto Navigli
Differentiable Multi-Agent Actor-Critic for Multi-Step Radiology Report Summarization	Sanjeev Kumar Karn, Ning Liu, Hinrich Schuetze, Oladimeji Farri
Direct Speech-to-Speech Translation With Discrete Units	Ann Lee, Peng-Jen Chen, Changhan Wang, Jiatao Gu, Sravya Popuri, Xutai Ma, Adam Polyak, Yossi Adi, Qing He, Yun Tang, Juan Pino, Wei-Ning Hsu
Discrete Opinion Tree Induction for Aspect-based Sentiment Analysis	Chenhua Chen, Zhiyang Teng, Zhongqing Wang, Yue Zhang
Discriminative Marginalized Probabilistic Neural Method for Multi-Document Summarization of Medical Literature	Gianluca Moro, Luca Ragazzi, Lorenzo Valgimigli, Davide Freddi
Disentangled Sequence to Sequence Learning for Compositional Generalization	Hao Zheng, Mirella Lapata
Distantly Supervised Named Entity Recognition via Confidence-Based Multi-Class Positive and Unlabeled Learning	Kang Zhou, Yuepei Li, Qi Li
Distributionally Robust Finetuning BERT for Covariate Drift in Spoken Language Understanding	Samuel Broscheit, Quynh Do, Judith Gaspers
Ditch the Gold Standard: Re-evaluating Conversational Question Answering	Huihan Li, Tianyu Gao, Manan Goenka, Danqi Chen
Divide and Denoise: Learning from Noisy Labels in Fine-Grained Entity Typing with Cluster-Wise Loss Correction	Kunyuan Pang, Haoyu Zhang, Jie Zhou, Ting Wang

Divide and Rule: Effective Pre-Training for Context-Aware Multi-Encoder Translation Models	Lorenzo Lupo, Marco Dinarelli, laurent besacier
Do self-supervised speech models develop human-like perception biases?	Juliette MILLET, Ewan Dunbar
Do Transformer Models Show Similar Attention Patterns to Task-Specific Human Gaze?	Oliver Eberle, Stephanie Brandl, Jonas Pilot, Anders Søgaard
DoCoGen: Domain Counterfactual Generation for Low Resource Domain Adaptation	Nitay Calderon, Eyal Ben-David, Amir Feder, Roi Reichart
Doctor Recommendation in Online Health Forums via Expertise Learning	Xiaoxin Lu, Yubo Zhang, Jing Li, Shi Zong
Does Recommend-Revise Produce Reliable Annotations? An Analysis on Missing Instances in DocRED	Quzhe Huang, Shibo Hao, Yuan Ye, Shengqi Zhu, Yansong Feng, Dongyan Zhao
Domain Adaptation in Multilingual and Multi-Domain Monolingual Settings for Complex Word Identification	George-Eduard Zaharia, Răzvan-Alexandru Smădu, Dumitru Clementin Cercel, Mihai Dascalu
Domain Knowledge Transferring for Pre-trained Language Model via Calibrated Activation Boundary Distillation	Dongha Choi, HongSeok Choi, Hyunju Lee
Down and Across: Introducing Crossword-Solving as a New NLP Benchmark	Saurabh Kulshreshtha, Olga Kovaleva, Namrata Shivagunde, Anna Rumshisky
DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization	Ziming Mao, Chen Henry Wu, Ansong Ni, Yusen Zhang, Rui Zhang, Tao Yu, Budhaditya Deb, Chenguang Zhu, Ahmed Hassan Awadallah, Dragomir Radev
Dynamic Global Memory for Document-level Argument Extraction	Xinya Du, Sha Li, Heng Ji
Dynamic Prefix-Tuning for Generative Template-based Event Extraction	Xiao Liu, Heyan Huang, Ge Shi, Bo Wang
Dynamic Schema Graph Fusion Network for Multi-Domain Dialogue State Tracking	Yue Feng, Aldo Lipani, Fanghua Ye, Qiang Zhang, Emine Yilmaz
e-CARE: a New Dataset for Exploring Explainable Causal Reasoning	Li Du, Xiao Ding, Kai Xiong, Ting Liu, Bing Qin
E-LANG: Energy-Based Joint Inferencing of Super and Swift Language Models	Mohammad Akbari, Amin Banitalebi-Dehkordi, Yong Zhang
EAG: Extract and Generate Multi-way Aligned Corpus for Complete Multilingual Neural Machine Translation	Yulin Xu, Zhen Yang, Fandong Meng, Jie Zhou
Early Stopping Based on Unlabeled Samples in Text Classification	HongSeok Choi, Dongha Choi, Hyunju Lee
Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-centric Summarization	Zhenjie Zhao, Yufang Hou, Dakuo Wang, Mo Yu, Chengzhong LIU, Xiaojuan Ma
Effective Token Graph Modeling using a Novel Labeling Strategy for Structured Sentiment Analysis	Wenxuan Shi, Fei Li, Jingye Li, Hao Fei, Donghong Ji
Efficient Cluster-Based k-Nearest-Neighbor Machine Translation	Dexin Wang, Kai Fan, Boxing Chen, Deyi Xiong
Efficient Hyper-parameter Search for Knowledge Graph Embedding	Yongqi Zhang, Zhanke Zhou, quanming yao, Yong Li

Efficient Unsupervised Sentence Compression by Fine-tuning Transformers with Reinforcement Learning	Demian Gholipour Ghalandari, Chris Hokamp, Georgiana Ifrim
End-to-End Modeling via Information Tree for One-Shot Natural Language Spatial Video Grounding	Mengze Li, Tianbao Wang, Haoyu Zhang, Shengyu Zhang, Zhou Zhao, Jiaxu Miao, Wenqiao Zhang, Wenming Tan, Jin Wang, PENG WANG, Shiliang Pu, Fei Wu
Enhanced Multi-Channel Graph Convolutional Network for Aspect Sentiment Triplet Extraction	Hao Chen, Zepeng Zhai, Fangxiang Feng, Ruifan Li, Xiaojie Wang
Enhancing Chinese Pre-trained Language Model via Heterogeneous Linguistics Graph	Yanzeng Li, Jiangxia Cao, Xin Cong, Zhenyu Zhang, Bowen Yu, Hongsong Zhu, Tingwen Liu
Enhancing Cross-lingual Natural Language Inference by Prompt-learning from Cross-lingual Templates	Kunxun Qi, Hai Wan, Jianfeng Du, Haolan Chen
Ensembling and Knowledge Distilling of Large Sequence Taggers for Grammatical Error Correction	Maksym Tarnavskiy, Artem Chernodub, Kostiantyn Omelianchuk
Entailment Graph Learning with Textual Entailment and Soft Transitivity	Zhibin Chen, Yansong Feng, Dongyan Zhao
Entity-based Neural Local Coherence Modeling	Sungho Jeon, Michael Strube
EntSUM: A Data Set for Entity-Centric Extractive Summarization	Mounica Maddela, Mayank Kulkarni, Daniel Preotiuc-Pietro
ePIC: Employing Proverbs in Context as a Benchmark for Abstract Language Understanding	Sayan Ghosh, Shashank Srivastava
EPT-X: An Expression-Pointer Transformer model that generates eXplanations for numbers	Bugeun Kim, Kyung Seo Ki, Sangkyu Rhim, gahgene gweon
Ethics Sheets for AI Tasks	Saif M. Mohammad
Evaluating Extreme Hierarchical Multi-label Classification	Enrique Amigo, Agustín D. Delgado
Evaluating Factuality in Text Simplification	Ashwin Devaraj, William Berkeley Sheffield, Byron C Wallace, Junyi Jessie Li
Expanding Pretrained Models to Thousands More Languages via Lexicon-based Adaptation	Xinyi Wang, Sebastian Ruder, Graham Neubig
Explanation Graph Generation via Pre-trained Language Models: An Empirical Study with Contrastive Learning	Swarnadeep Saha, Prateek Yadav, Mohit Bansal
Exploring and Adapting Chinese GPT to Pinyin Input Method	Minghuan Tan, Yong Dai, Duyu Tang, Zhangyin Feng, Guoping Huang, Jing Jiang, Jiwei Li, Shuming Shi
ExtEnD: Extractive Entity Disambiguation	Edoardo Barba, Luigi Procopio, Roberto Navigli
Fair and Argumentative Language Modeling for Computational Argumentation	Carolin Holtermann, Anne Lauscher, Simone Paolo Ponzetto
FairLex: A Multilingual Benchmark for Evaluating Fairness in Legal Text Processing	Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Felix Schwemer, Anders Søgaard

FaiRR: Faithful and Robust Deductive Reasoning over Natural Language	Soumya Sanyal, Harman Singh, Xiang Ren
Faithful or Extractive? On Mitigating the Faithfulness-Abstractiveness Trade-off in Abstractive Summarization	Faisal Ladhak, Esin Durmus, He He, Claire Cardie, Kathleen McKeown
Fantastic Questions and Where to Find Them: FairytaleQA -- An Authentic Dataset for Narrative Comprehension	Ying Xu, Dakuo Wang, Mo Yu, Daniel Ritchie, Bingsheng Yao, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Nora Bradford, Branda Sun, Tran Bao Hoang, Yisi Sang, Yufang Hou, Xiaojuan Ma, Diyi Yang, Nanyun Peng, Zhou Yu, Mark Warschauer
Fantastically Ordered Prompts and Where to Find Them: Overcoming Few-Shot Prompt Order Sensitivity	Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, Pontus Stenetorp
FaVIQ: FAct Verification from Information-seeking Questions	Jungsoo Park, Sewon Min, Jaewoo Kang, Luke Zettlemoyer, Hannaneh Hajishirzi
Feeding What You Need by Understanding What You Learned	Xiaoqiang Wang, Bang Liu, Fangli Xu, Bo Long, Siliang Tang, Lingfei Wu
Few-Shot Class-Incremental Learning for Named Entity Recognition	Rui Wang, Tong Yu, Handong Zhao, Sungchul Kim, Subrata Mitra, Ruiyi Zhang, Ricardo Henao
Few-shot Controllable Style Transfer for Low-Resource Multilingual Settings	Kalpesh Krishna, Deepak Nathani, Xavier Garcia, Bidisha Samanta, Partha Talukdar
Few-Shot Learning with Siamese Networks and Label Tuning	Thomas Müller, Guillermo Pérez-Torró, Marc Franco-Salvador
Few-shot Named Entity Recognition with Self-describing Networks	Jiawei Chen, Qing Liu, Hongyu Lin, Xianpei Han, Le Sun
Few-Shot Tabular Data Enrichment Using Fine-Tuned Transformer Architectures	Asaf Harari, Gilad Katz
FewNLU: Benchmarking State-of-the-Art Methods for Few-Shot Natural Language Understanding	Yanan Zheng, Jing Zhou, Yujie Qian, Ming Ding, Chonghua Liao, Jian Li, Ruslan Salakhutdinov, Jie Tang, Sebastian Ruder, Zhilin Yang
FIBER: Fill-in-the-Blanks as a Challenging Video Understanding Evaluation Framework	Santiago Castro, Ruoyao Wang, Pingxuan Huang, Ian Stewart, Oana Ignat, Nan Liu, Jonathan C. Stroud, Rada Mihalcea
Finding Structural Knowledge in Multimodal-BERT	Victor Siemen Janusz Milewski, Miryam de Lhoneux, Marie-Francine Moens
Fine- and Coarse-Granularity Hybrid Self-Attention for Efficient BERT	Jing Zhao, Yifan Wang, Junwei Bao, Youzheng Wu, Xiaodong He
Fine-Grained Controllable Text Generation Using Non-Residual Prompting	Fredrik Carlsson, Joey Öhman, Fangyu Liu, Severine Verlinden, Joakim Nivre, Magnus Sahlgren
FINER: Financial Numeric Entity Recognition for XBRL Tagging	Lefteris Loukas, Manos Fergadiotis, Ilias Chalkidis, Eirini Spyropoulou, Prodromos Malakasiotis, Ion Androutsopoulos, Georgios Paliouras

Flexible Generation from Fragmentary Linguistic Input	Peng Qian, Roger P. Levy
FlipDA: Effective and Robust Data Augmentation for Few-Shot Learning	Jing Zhou, Yanan Zheng, Jie Tang, Jian Li, Zhilin Yang
Flooding-X: Improving BERT's Resistance to Adversarial Attacks via Loss-Restricted Fine-Tuning	Qin Liu, Rui Zheng, Bao Rong, Jingyi Liu, ZhiHua Liu, Zhanzhan Cheng, Liang Qiao, Tao Gui, Qi Zhang, Xuanjing Huang
Flow-Adapter Architecture for Unsupervised Machine Translation	Yihong Liu, Haris Jabbar, Hinrich Schuetze
FormNet: Structural Encoding beyond Sequential Modeling in Form Document Information Extraction	Chen-Yu Lee, Chun-Liang Li, Timothy Dozat, Vincent Perot, Guolong Su, Nan Hua, Joshua Ainslie, Renshen Wang, Yasuhisa Fujii, Tomas Pfister
FORTAP: Using Formulas for Numerical-Reasoning-Aware Table Pretraining	Zhoujun Cheng, Haoyu Dong, Ran Jia, pengfei wu, Shi Han, Fan Cheng, Dongmei Zhang
French CrowS-Pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English	Aur�lie N�v�ol, Yoann Dupont, Julien Bezan�on, Kar�n Fort
From Simultaneous to Streaming Machine Translation by Leveraging Streaming History	Javier Iranzo Sanchez, Jorge Civera, Alfons Juan-C�scar
From text to talk: Harnessing conversational corpora for humane and diversity-aware language technology	Mark Dingemanse, Andreas Liesenfeld
From the Detection of Toxic Spans in Online Discussions to the Analysis of Toxic-to-Civil Transfer	John Pavlopoulos, Leo Laugier, ALEXANDROS XENOS, Jeffrey S. Sorensen, Ion Androutsopoulos
FrugalScore: Learning Cheaper, Lighter and Faster Evaluation Metrics for Automatic Text Generation	Moussa KAMAL EDDINE, Guokan Shang, Antoine J.-P. Tixier, Michalis Vazirgiannis
Fully Hyperbolic Neural Networks	Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou
Fully-Semantic Parsing and Generation: the BabelNet Meaning Representation	Abelardo Carlos Martinez Lorenzo, Marco Maru, Roberto Navigli
Generated Knowledge Prompting for Commonsense Reasoning	Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, Hannaneh Hajishirzi
Generating Biographies on Wikipedia: The Impact of Gender Bias on the Retrieval-Based Generation of Women Biographies	Angela Fan, Claire Gardent
Generating Data to Mitigate Spurious Correlations in Natural Language Inference Datasets	Yuxiang Wu, Matt Gardner, Pontus Stenetorp, Pradeep Dasigi
Generating Scientific Claims for Zero-Shot Scientific Fact Checking	Dustin Wright, David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Isabelle Augenstein, Lucy Lu Wang
Generating Scientific Definitions with Controllable Complexity	Tal August, Katharina Reinecke, Noah Smith
Generative Pretraining for Paraphrase Evaluation	Jack Weston, Raphael Lenain, Udeepa Meepegama, Emil Fristed

GL-CLeF: A Global--Local Contrastive Learning Framework for Cross-lingual Spoken Language Understanding	Libo Qin, Qiguang Chen, Tianbao Xie, Qixin Li, Jian-Guang Lou, Wanxiang Che, Min-Yen Kan
GLM: General Language Model Pretraining with Autoregressive Blank Infilling	Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, Jie Tang
GlobalWoZ: Globalizing MultiWoZ to Develop Multilingual Task-Oriented Dialogue Systems	Bosheng Ding, Junjie Hu, Lidong Bing, Mahani Aljunied, Shafiq Joty, Luo Si, Chunyan Miao
Good Examples Make A Faster Learner: Simple Demonstration-based Learning for Low-resource NER	Dong-Ho Lee, Akshen Kadakia, Kangmin Tan, Mahak Agarwal, Xinyu Feng, Takashi Shibuya, Ryosuke Mitani, Toshiyuki Sekiya, Jay Pujara, Xiang Ren
GPT-D: Inducing Dementia-related Linguistic Anomalies by Deliberate Degradation of Artificial Neural Language Models	Changye Li, David Knopman, Weizhe Xu, Trevor Cohen, Serguei V. S. Pakhomov
Graph Enhanced Contrastive Learning for Radiology Findings Summarization	Jinpeng Hu, Zhuo Li, Zhihong Chen, Zhen Li, Xiang Wan, Tsung-Hui Chang
Graph Pre-training for AMR Parsing and Generation	Xuefeng Bai, Yulong Chen, Yue Zhang
Guided Attention Multimodal Multitask Financial Forecasting with Inter-Company Relationships and Global and Local News	Gary Ang, Ee-Peng Lim
Hallucinated but Factual! Inspecting the Factuality of Hallucinations in Abstractive Summarization	Meng Cao, Yue Dong, Jackie CK Cheung
Headed-Span-Based Projective Dependency Parsing	Songlin Yang, Kewei Tu
HeterMPC: A Heterogeneous Graph Neural Network for Response Generation in Multi-Party Conversations	Jia-Chen Gu, Chao-Hong Tan, Chongyang Tao, Zhen-Hua Ling, Huang Hu, Xiubo Geng, Daxin Jiang
HIBRIDS: Attention with Hierarchical Biases for Structure-aware Long Document Summarization	Shuyang Cao, Lu Wang
Hierarchical Sketch Induction for Paraphrase Generation	Tom Hosking, Hao Tang, Mirella Lapata
HiTab: A Hierarchical Table Dataset for Question Answering and Natural Language Generation	Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, Dongmei Zhang
HOLM: Hallucinating Objects with Language Models for Referring Expression Recognition in Partially-Observed Scenes	Volkan Cirik, Louis-Philippe Morency, Taylor Berg-Kirkpatrick
How can NLP Help Revitalize Endangered Languages? A Case Study and Roadmap for the Cherokee Language	Shiyue Zhang, Ben Frey, Mohit Bansal
How Do Seq2Seq Models Perform on End-to-End Data-to-Text Generation?	Xunjian Yin, Xiaojun Wan
How Do We Answer Complex Questions: Discourse Structure of Long-form Answers	Fangyuan Xu, Junyi Jessy Li, Eunsol Choi
Human Evaluation and Correlation with Automatic Metrics in Consultation Note Generation	Francesco Moramarco, Alex Papadopoulos Korfiatis, Mark Perera, Damir JURIC, Jack Flann, Ehud Reiter, Anya Belz, Aleksandar Savkov

Hybrid Semantics for Goal-Directed Natural Language Generation	Connor Baumler, Soumya Ray
Hypergraph Transformer: Weakly-Supervised Multi-hop Reasoning for Knowledge-based Visual Question Answering	Yu-Jung Heo, Eun-Sol Kim, Woo Suk Choi, Byoung-Tak Zhang
Hyperlink-induced Pre-training for Passage Retrieval in Open-domain Question Answering	Jiawei Zhou, Xiaoguang Li, Lifeng Shang, Lan Luo, Ke Zhan, Enrui Hu, Xinyu Zhang, Hao Jiang, Zhao Cao, Fan Yu, Xin Jiang, Qun Liu, Lei Chen
IAM: A Comprehensive and Large-Scale Dataset for Integrated Argument Mining Tasks	Liyang Cheng, Lidong Bing, Ruidan He, Qian Yu, Yan Zhang, Luo Si
Identifying Chinese Opinion Expressions with Extremely-Noisy Crowdsourcing Annotations	Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, Xiaobin Wang, Min Zhang
Identifying Moments of Change from Longitudinal User Text	Adam Tsakalidis, Federico Nanni, Anthony Hills, Jenny Chim, Jiayu Song, Maria Liakata
Identifying the Human Values behind Arguments	Johannes Kiesel, Milad Alshomary, Nicolas Handke, Xiaoni Cai, Henning Wachsmuth, Benno Stein
ILDAE: Instance-Level Difficulty Analysis of Evaluation Data	Neeraj Varshney, Swaroop Mishra, Chitta Baral
Image Retrieval from Contextual Descriptions	Benno Krojer, Vaibhav Adlakha, Vibhav Vineet, Yash Goyal, Edoardo Ponti, Siva Reddy
Impact of Evaluation Methodologies on Code Summarization	Pengyu Nie, Jiyang Zhang, Junyi Jessy Li, Ray Mooney, Milos Gligoric
IMPLI: Investigating NLI Models' Performance on Figurative Language	Kevin Stowe, Prasetya Ajie Utama, Iryna Gurevych
Improving Compositional Generalization with Self-Training for Data-to-Text Generation	Sanket Vaibhav Mehta, Jinfeng Rao, Yi Tay, Mihir Kale, Ankur P Parikh, Emma Strubell
Improving Event Representation via Simultaneous Weakly Supervised Contrastive Learning and Clustering	Jun Gao, Wei Wang, Changlong Yu, Huan Zhao, Wilfred Ng, Ruifeng Xu
Improving Generalizability in Implicitly Abusive Language Detection with Concept Activation Vectors	Isar Nejadgholi, Kathleen C. Fraser, Svetlana Kiritchenko
Improving Machine Reading Comprehension with Contextualized Commonsense Knowledge	Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Claire Cardie
Improving Meta-learning for Low-resource Text Classification and Generation via Memory Imitation	Yingxiu Zhao, Zhiliang Tian, Huaxiu Yao, Yinhe Zheng, Dongkyu Lee, YIPING SONG, Jian SUN, Nevin Zhang
Improving Multi-label Malevolence Detection in Dialogues through Multi-faceted Label Correlation Enhancement	Yangjun Zhang, Pengjie Ren, Wentao Deng, Zhumin Chen, Maarten de Rijke
Improving Personalized Explanation Generation through Visualization	Shijie Geng, Zuohui Fu, Yingqiang Ge, Lei Li, Gerard de Melo, Yongfeng Zhang
Improving the Generalizability of Depression Detection by Leveraging Clinical Questionnaires	Thong Nguyen, Andrew Yates, Ayah Zirikly, Bart Desmet, Arman Cohan
Improving Time Sensitivity for Question Answering over Temporal	Chao Shang, Guangtao Wang, Peng Qi, Jing

Knowledge Graphs	Huang
Improving Word Translation via Two-Stage Contrastive Learning	Yaoyiran Li, Fangyu Liu, Nigel Collier, Anna Korhonen, Ivan Vulić
Imputing Out-of-Vocabulary Embeddings with LOVE Makes Language Models Robust with Little Cost	Lihu Chen, Gael Varoquaux, Fabian M. Suchanek
Incorporating Hierarchy into Text Encoder: a Contrastive Learning Approach for Hierarchical Text Classification	Zihan Wang, Peiyi Wang, Lianzhe Huang, Xin Sun, Houfeng Wang
Incorporating Stock Market Signals for Twitter Stance Detection	Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd, Nigel Collier
Inducing Positive Perspectives with Text Reframing	Caleb Ziems, Minzhi Li, Anthony Zhang, Diyi Yang
Inferring Rewards from Language in Context	Jessy Lin, Daniel Fried, Dan Klein, Anca Dragan
infy-former: Infinite Memory Transformer	Pedro Henrique Martins, Zita Marinho, Andre Martins
Integrating Vectorized Lexical Constraints for Neural Machine Translation	Shuo Wang, Zhixing Tan, Yang Liu
Interactive Word Completion for Plains Cree	William Abbott Lane, Atticus Galvin Harrigan, Antti Arppe
Internet-Augmented Dialogue Generation	Mojtaba Komeili, Kurt Shuster, Jason E Weston
Interpretability for Language Learners Using Example-Based Grammatical Error Correction	Masahiro Kaneko, Sho Takase, Ayana Niwa, Naoaki Okazaki
Interpreting Character Embeddings With Perceptual Representations: The Case of Shape, Sound, and Color	Sidsel Boldsen, Manex Agirrezabal, Nora Hollenstein
Investigating Failures of Automatic Translation in the Case of Unambiguous Gender	Adi Renduchintala, Adina Williams
Investigating Non-local Features for Neural Constituency Parsing	Leyang Cui, Sen Yang, Yue Zhang
Is Attention Explanation? An Introduction to the Debate	Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François, Patrick Watrin
Is GPT-3 Text Indistinguishable from Human Text? Scarecrow: A Framework for Scrutinizing Machine Text	Yao Dou, Maxwell Forbes, Rik Koncel-Kedziorski, Noah Smith, Yejin Choi
It is AI's Turn to Ask Humans a Question: Question-Answer Pair Generation for Children's Story Books	Bingsheng Yao, Dakuo Wang, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Mo Yu, Ying Xu
JointCL: A Joint Contrastive Learning Framework for Zero-Shot Stance Detection	Bin Liang, Qinglin Zhu, Xiang Li, Min Yang, Lin Gui, Yulan He, Ruifeng Xu
Just Rank: Rethinking Evaluation with Word and Sentence Similarities	Bin Wang, C.-C. Jay Kuo, Haizhou Li
KaFSP: Knowledge-Aware Fuzzy Semantic Parsing for Conversational Question Answering over a Large-Scale Knowledge Base	Junzhuo Li, Deyi Xiong

KenMeSH: Knowledge-enhanced End-to-end Biomedical Text Labelling	Xindi Wang, Robert Mercer, Frank Rudzicz
Keywords and Instances: A Hierarchical Contrastive Learning Framework Unifying Hybrid Granularities for Text Generation	Mingzhe Li, XieXiong Lin, Xiuying Chen, Jinxiong Chang, Qishen Zhang, Feng Wang, Taifeng Wang, Zhongyi Liu, Wei Chu, Dongyan Zhao, Rui Yan
KG-FiD: Infusing Knowledge Graph in Fusion-in-Decoder for Open-Domain Question Answering	Donghan Yu, Chenguang Zhu, Yuwei Fang, Wenhao Yu, Shuohang Wang, Yichong Xu, Xiang Ren, Yiming Yang, Michael Zeng
KinyaBERT: a Morphology-aware Kinyarwanda Language Model	Antoine Nzeyimana, Andre Niyongabo Rubungo
KNN-Contrastive Learning for Out-of-Domain Intent Classification	Yunhua Zhou, Peiju Liu, Xipeng Qiu
Knowledge Enhanced Reflection Generation for Counseling Dialogues	Siqi Shen, Veronica Perez-Rosas, Charles Welch, Soujanya Poria, Rada Mihalcea
Knowledge Neurons in Pretrained Transformers	Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, Furu Wei
Knowledgeable Prompt-tuning: Incorporating Knowledge into Prompt Verbalizer for Text Classification	Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, Maosong Sun
KQA Pro: A Dataset with Explicit Compositional Programs for Complex Question Answering over Knowledge Base	Shulin Cao, Jiaxin Shi, Liangming Pan, Lunyiu Nie, Yutong Xiang, Lei Hou, Juanzi Li, Bin He, Hanwang Zhang
Label Semantic Aware Pre-training for Few-shot Text Classification	Aaron Mueller, Jason Krone, Salvatore Romeo, Saab Mansour, Elman Mansimov, Yi Zhang, Dan Roth
LAGr: Label Aligned Graphs for Better Systematic Generalization in Semantic Parsing	Dora Jambor, Dzmitry Bahdanau
Language-agnostic BERT Sentence Embedding	Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, Wei Wang
Language-Agnostic Meta-Learning for Low-Resource Text-to-Speech with Articulatory Features	Florian Lux, Thang Vu
Large Scale Substitution-based Word Sense Induction	Matan Eyal, Shoval Sadde, Hillel Taub-Tabib, Yoav Goldberg
latent-GLAT: Glancing at Latent Variables for Parallel Text Generation	Yu Bao, Hao Zhou, Shujian Huang, dongqi wang, Lihua Qian, Xinyu Dai, Jiajun Chen, Lei Li
Learn to Adapt for Generalized Zero-Shot Text Classification	Yiwen Zhang, Caixia Yuan, Xiaojie Wang, Ziwei Bai, Yongbin Liu
Learned Incremental Representations for Parsing	Nikita Kitaev, Thomas Lu, Dan Klein
Learning Adaptive Segmentation Policy for End-to-End Simultaneous Translation	Ruiqing Zhang, Zhongjun He, Hua Wu, Haifeng Wang
Learning Confidence for Transformer-based Neural Machine Translation	Yu Lu, Jiali Zeng, Jiajun Zhang, Shuangzhi Wu, Mu Li

Learning Disentangled Representations of Negation and Uncertainty	Jake A Vasilakes, Chrysoula Zerva, Makoto Miwa, Sophia Ananiadou
Learning Disentangled Semantic Representations for Zero-Shot Cross-Lingual Transfer in Multilingual Machine Reading Comprehension	Linjuan Wu, Shaojuan Wu, Xiaowang Zhang, Deyi Xiong, Shizhan Chen, Zhiqiang Zhuang, Zhiyong Feng
Learning Disentangled Textual Representations via Statistical Measures of Similarity	Pierre Colombo, Guillaume Staerman, Nathan Noiry, Pablo Piantanida
Learning From Failure: Data Capture in an Australian Aboriginal Community	Eric Le Ferrand, Steven Bird, Laurent BesacierUGA
Learning from Sibling Mentions with Scalable Graph Inference in Fine-Grained Entity Typing	Yi Chen, Jiayang Cheng, Haiyun Jiang, Lemao Liu, Haisong Zhang, Shuming Shi, Ruifeng Xu
Learning Functional Distributional Semantics with Visual Data	Yinhong Liu, Guy Emerson
Learning Non-Autoregressive Models from Search for Unsupervised Sentence Summarization	Puyuan Liu, Chenyang Huang, Lili Mou
Learning the Beauty in Songs: Neural Singing Voice Beautifier	Jinglin Liu, Chengxi Li, Yi Ren, Zhiying Zhu, Zhou Zhao
Learning to Generalize to More: Continuous Semantic Augmentation for Neural Machine Translation	Xiangpeng Wei, Heng Yu, Yue Hu, Rongxiang Weng, Weihua Luo, Rong Jin
Learning to Generate Programs for Table Fact Verification via Structure-Aware Semantic Parsing	Suixin Ou, Yongmei Liu
Learning to Imagine: Integrating Counterfactual Thinking in Neural Discrete Reasoning	Moxin Li, Fuli Feng, Hanwang Zhang, Xiangnan He, Fengbin ZHU, Tat-Seng Chua
Learning to Mediate Disparities Towards Pragmatic Communication	Yuwei Bao, Sayan Ghosh, Joyce Chai
Learning to Rank Visual Stories From Human Ranking Data	Chi-Yang Hsu, Yun-Wei Chu, Vincent Chen, Kuan-Chieh Lo, Chacha Chen, Ting-Hao Huang, Lun-Wei Ku
Learning to Reason Deductively: Math Word Problem Solving as Complex Relation Extraction	Zhanming Jie, Jierui Li, Wei Lu
Learning When to Translate for Streaming Speech	Qian qian Dong, Yaoming Zhu, Mingxuan Wang, Lei Li
Legal Judgment Prediction via Event Extraction with Constraints	Yi Feng, Chuanyi Li, Vincent Ng
Length Control in Abstractive Summarization by Pretraining Information Selection	Yizhu Liu, Qi Jia, Kenny Q. Zhu
Letters From the Past: Modeling Historical Sound Change Through Diachronic Character Embeddings	Sidsel Boldsen, Patrizia Paggio
Leveraging Relaxed Equilibrium by Lazy Transition for Sequence Modeling	Xi Ai, Bin Fang
Leveraging Similar Users for Personalized Language Modeling with Limited Data	Charles Welch, Chenxi Gu, Jonathan K Kummerfeld, Veronica Perez-Rosas, Rada Mihalcea

Leveraging Task Transferability to Meta-learning for Clinical Section Classification with Limited Data	Zhuohao Chen, Jangwon Kim, Ram Bhakta, MUSTAFA Y SIR
Leveraging Unimodal Self-Supervised Learning for Multimodal Audio-Visual Speech Recognition	Xichen Pan, Peiyu Chen, Yichen Gong, Helong Zhou, Xinbing Wang, Zhouhan Lin
Leveraging Visual Knowledge in Language Tasks: An Empirical Study on Intermediate Pre-training for Cross-Modal Knowledge Transfer	Woojeong Jin, Dong-Ho Lee, Chenguang Zhu, Jay Pujara, Xiang Ren
Leveraging Wikipedia article evolution for promotional tone detection	Christine De Kock, Andreas Vlachos
LexGLUE: A Benchmark Dataset for Legal Language Understanding in English	Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, Nikolaos Aletras
Lexical Knowledge Internalization for Neural Dialog Generation	Zhiyong Wu, Wei Bi, Xiang Li, Lingpeng Kong, Ben Kao
LexSubCon: Integrating Knowledge from Lexical Resources into Contextual Embeddings for Lexical Substitution	George Michalopoulos, Ian McKillop, Alexander Wong, Helen Hong Chen
Life after BERT: What do Other Muppets Understand about Language?	Vladislav Lialin, Kevin Zhao, Namrata Shivagunde, Anna Rumshisky
LiLT: A Simple yet Effective Language-Independent Layout Transformer for Structured Document Understanding	Jiapeng Wang, Lianwen Jin, Kai Ding
LinkBERT: Pretraining Language Models with Document Links	Michihiro Yasunaga, Jure Leskovec, Percy Liang
Lite Unified Modeling for Discriminative Reading Comprehension	Yilin Zhao, hai zhao, Libin Shen, Yinggong Zhao
Local Languages, Third Spaces, and other High-Resource Scenarios	Steven Bird
Logic Traps in Evaluating Attribution Scores	Yiming Ju, Yuanzhe Zhang, Zhao Yang, Zhongtao Jiang, Kang Liu, Jun Zhao
Long-range Sequence Modeling with Predictable Sparse Attention	Yimeng Zhuang, Jing Zhang, Mei Tu
Low-Rank Softmax Can Have Unargmaxable Classes in Theory but Rarely in Practice	Andreas Grivas, Nikolay Bogoychev, Adam Lopez
M3ED: Multi-modal Multi-scene Multi-label Emotional Dialogue Database	Jinming Zhao, Teggan Zhang, Jingwen Hu, Yuchen Liu, Qin Jin, Xinchao Wang, Haizhou Li
Make the Best of Cross-lingual Transfer: Evidence from POS Tagging with over 100 Languages	Wietse de Vries, Martijn Wieling, Malvina Nissim
Making Transformers Solve Compositional Tasks	Santiago Ontanon, Joshua Ainslie, Zachary Fisher, Vaclav Cvicek
MarkupLM: Pre-training of Text and Markup Language for Visually Rich Document Understanding	Junlong Li, Yiheng Xu, Lei Cui, Furu Wei
Match the Script, Adapt if Multilingual: Analyzing the Effect of Multilingual Pretraining on Cross-lingual Transferability	Yoshinari Fujinuma, Jordan Lee Boyd-Graber, Katharina Kann
Measuring and Mitigating Name Biases in Neural Machine Translation	Jun Wang, Benjamin I. P. Rubinstein, Trevor Cohn

Measuring Fairness of Text Classifiers via Prediction Sensitivity	Satyapriya Krishna, Rahul Gupta, Apurv Verma, Jwala Dhamala, Yada Pruksachatkun, Kai-Wei Chang
MELM: Data Augmentation with Masked Entity Language Modeling for Low-Resource NER	Ran Zhou, Xin Li, Ruidan He, Lidong Bing, Erik Cambria, Luo Si, Chunyan Miao
Memorisation versus Generalisation in Pre-trained Language Models	Michael Tänzler, Sebastian Ruder, Marek Rei
MemSum: Extractive Summarization of Long Documents Using Multi-Step Episodic Markov Decision Processes	Nianlong Gu, Elliott Ash, Richard Hahnloser
Meta-Learning for Fast Cross-Lingual Adaptation in Dependency Parsing	Anna Langedijk, Verna Dankers, Phillip Lippe, Sander Emile Bos, Bryan Cardenas Guevara, Helen Yannakoudakis, Ekaterina Shutova
Meta-learning via Language Model In-context Tuning	Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, He He
Metaphors in Pre-Trained Language Models: Probing and Generalization Across Datasets and Languages	Ehsan Aghazadeh, Mohsen Fayyaz, Yadollah Yaghoobzadeh
MILIE: Modular & Iterative Multilingual Open Information Extraction	Bhushan Kotnis, Kiril Gashteovski, Daniel Onoro Rubio, Ammar Shaker, Vanesa Rodriguez-Tembras, Makoto Takamoto, Mathias Niepert, Carolin Lawrence
MINER: Improving Out-of-Vocabulary Named Entity Recognition from an Information Theoretic Perspective	Xiao Wang, Shihan Dou, Limao Xiong, Yicheng Zou, Qi Zhang, Tao Gui, Liang Qiao, Zhanzhan Cheng, Xuanjing Huang
MISC: A Mixed Strategy-Aware Model integrating COMET for Emotional Support Conversation	Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, Rui Yan
Misinfo Reaction Frames: Reasoning about Readers' Reactions to News Headlines	Saadia Gabriel, Skyler Hallinan, Maarten Sap, Pemi Nguyen, Franziska Roesner, Eunsol Choi, Yejin Choi
Mix and Match: Learning-free Controllable Text Generation using Energy Language Models	Fatemehsadat Mireshghallah, Kartik Goyal, Taylor Berg-Kirkpatrick
mLUKE: The Power of Entity Representations in Multilingual Pretrained Language Models	Ryokan Ri, Ikuya Yamada, Yoshimasa Tsuruoka
MMCoQA: Conversational Question Answering over Text, Tables, and Images	Yongqi Li, Wenjie Li, Liqiang Nie
Modeling Dual Read/Write Paths for Simultaneous Machine Translation	Shaolei Zhang, Yang Feng
Modeling Hierarchical Syntax Structure with Triplet Position for Source Code Summarization	Juncai Guo, Jin Liu, Yao Wan, Li Li, Pingyi Zhou
Modeling Multi-hop Question Answering as Single Sequence Prediction	Semih Yavuz, Kazuma Hashimoto, Yingbo Zhou, Nitish Shirish Keskar, Caiming Xiong
Modeling Persuasive Discourse to Adaptively Support Students' Argumentative Writing	Thiemo Wambsganss, Christina Niklaus
Modeling Syntactic-Semantic Dependency Correlations in Semantic Role	Junjie Chen, Xiangheng He, Yusuke Miyao

Labeling Using Mixture Models	
Modeling Temporal-Modal Entity Graph for Procedural Multimodal Machine Comprehension	Huibin Zhang, Zhengkun Zhang, YAO ZHANG, Jun Wang, Yufan Li, ning jiang, xin wei, Zhenglu Yang
Modeling U.S. State-Level Policies by Extracting Winners and Losers from Legislative Texts	Maryam Davoodi, Eric Waltenburg, Dan Goldwasser
MPII: Multi-Level Mutual Promotion for Inference and Interpretation	Yan Liu, Sanyuan Chen, Yazheng Yang, Qi Dai
MSCTD: A Multimodal Sentiment Chat Translation Dataset	Yunlong Liang, Fandong Meng, Jinan Xu, Yufeng Chen, Jie Zhou
MSP: Multi-Stage Prompting for Making Pre-trained Language Models Better Translators	Zhixing Tan, Xiangwen Zhang, Shuo Wang, Yang Liu
Multi Task Learning For Zero Shot Performance Prediction of Multilingual Models	Kabir Ahuja, Shanu Kumar, Sandipan Dandapat, Monojit Choudhury
Multi-Granularity Structural Knowledge Distillation for Language Model Compression	Chang Liu, Chongyang Tao, Jiazhan Feng, Dongyan Zhao
Multi-Modal Sarcasm Detection via Cross-Modal Graph Convolutional Network	Bin Liang, Chenwei Lou, Xiang Li, Min Yang, Lin Gui, Yulan He, Wenjie Pei, Ruifeng Xu
Multi-Party Empathetic Dialogue Generation: A New Task for Dialog Systems	Ling.Yu Zhu, Zhengkun Zhang, Jun Wang, hongbin wang, haiying wu, Zhenglu Yang
Multi-Task Pre-Training for Plug-and-Play Task-Oriented Dialogue System	Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, Yi Zhang
Multi-View Document Representation Learning for Open-Domain Dense Retrieval	Shunyu Zhang, Yaobo Liang, MING GONG, Daxin Jiang, Nan Duan
MultiHiertt: Numerical Reasoning over Multi Hierarchical Tabular and Textual Data	Yilun Zhao, Yunxiang Li, Chenying Li, Rui Zhang
Multilingual Detection of Personal Employment Status on Twitter	Manuel Tonneau, Dhaval Adjudah, Joao Palotti, Nir Grinberg, Samuel Fraiberger
Multilingual Document-Level Translation Enables Zero-Shot Transfer From Sentences to Documents	Biao Zhang, Ankur Bapna, Melvin Johnson, Ali Dabirmoghaddam, Naveen Arivazhagan, Orhan Firat
Multilingual Generative Language Models for Zero-Shot Cross-Lingual Event Argument Extraction	Kuan-Hao Huang, I-Hung Hsu, Prem Natarajan, Kai-Wei Chang, Nanyun Peng
Multilingual Knowledge Graph Completion with Self-Supervised Adaptive Graph Alignment	Zijie Huang, zheng li, Haoming Jiang, Tianyu Cao, Hanqing Lu, Bing Yin, Karthik Subbian, Yizhou Sun, Wei Wang
Multilingual Mix: Example Interpolation Improves Multilingual Neural Machine Translation	Yong Cheng, Ankur Bapna, Orhan Firat, Yuan Cao, Pidong Wang, Wolfgang Macherey
Multilingual Molecular Representation Learning via Contrastive Pre-training	Zhihui Guo, Pramod Sharma, Andy Daniel Martinez, Liang Du, Robin Abraham
Multilingual unsupervised sequence segmentation transfers to extremely	C. M. Downey, Shannon Drizin, Levon

low-resource languages	Haroutunian, Shivin Thukral
Multimodal Dialogue Response Generation	Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng, Daxin Jiang
Multimodal fusion via cortical network inspired losses	Shiv Shankar
Multimodal Sarcasm Target Identification in Tweets	Jiquan Wang, Lin Sun, Yi Liu, Meizhi Shao, Zengwei Zheng
Multitasking Framework for Unsupervised Simple Definition Generation	Cunliang Kong, Yun Chen, Hengyuan Zhang, Liner Yang, Erhong Yang
Nested Named Entity Recognition as Latent Lexicalized Constituency Parsing	Chao Lou, Songlin Yang, Kewei Tu
Nested Named Entity Recognition with Span-level Graphs	Juncheng Wan, Dongyu Ru, Weinan Zhang, Yong Yu
Neural Label Search for Zero-Shot Multi-Lingual Extractive Summarization	Ruipeng Jia, Xingxing Zhang, Yanan Cao, Zheng Lin, Shi Wang, Furu Wei
Neural Machine Translation with Phrase-Level Universal Visual Representations	Qingkai Fang, Yang Feng
Neural Pipeline for Zero-Shot Data-to-Text Generation	Zdeněk Kasner, Ondrej Dusek
Neural reality of argument structure constructions	Bai Li, Zining Zhu, Guillaume Thomas, Frank Rudzicz, Yang Xu
New Intent Discovery with Pre-training and Contrastive Learning	Yuwei Zhang, Haode Zhang, Li-Ming Zhan, Xiao-Ming Wu, Albert Y.S. Lam
Nibbling at the Hard Core of Word Sense Disambiguation	Marco Maru, Simone Conia, Michele Bevilacqua, Roberto Navigli
Noisy Channel Language Model Prompting for Few-Shot Text Classification	Sewon Min, Mike Lewis, Hannaneh Hajishirzi, Luke Zettlemoyer
Non-neural Models Matter: a Re-evaluation of Neural Referring Expression Generation Systems	Fahime Same, Guanyi Chen, Kees Van Deemter
Not always about you: Prioritizing community needs when developing endangered language technology	Zoey Liu, Crystal Richardson, Richard John Hatcher, Emily Tucker Prud'hommeaux
NumGLUE: A Suite of Fundamental yet Challenging Mathematical Reasoning Tasks	Swaroop Mishra, Arindam Mitra, Neeraj Varshney, Bhavdeep Singh Sachdeva, Peter Clark, Chitta Baral, Ashwin Kalyan
ODE Transformer: An Ordinary Differential Equation-Inspired Model for Sequence Generation	Bei Li, Quan Du, Tao Zhou, Yi Jing, Shuhan Zhou, Xin Zeng, Tong Xiao, JingBo Zhu, Xuebo Liu, Min Zhang
OIE@OIA: an Adaptable and Efficient Open Information Extraction Framework	Xin Wang, Minlong Peng, Mingming Sun, Ping Li
On Continual Model Refinement in Out-of-Distribution Data Streams	Bill Yuchen Lin, Sida Wang, Xi Victoria Lin, Robin Jia, Lin Xiao, Xiang Ren, Scott Yih

On the Calibration of Pre-trained Language Models using Mixup Guided by Area Under the Margin and Saliency	Seo Yeon Park, Cornelia Caragea
On The Ingredients of an Effective Zero-shot Semantic Parser	Pengcheng Yin, John Frederick Wieting, Avirup Sil, Graham Neubig
On the Robustness of Offensive Language Classifiers	Jonathan Rusert, Zubair Shafiq, Padmini Srinivasan
On the Robustness of Question Rewriting Systems to Questions of Varying Hardness	Hai Ye, Hwee Tou Ng, Wenjuan Han
On the Sensitivity and Stability of Model Interpretations in NLP	Fan Yin, Zhouxing Shi, Cho-Jui Hsieh, Kai-Wei Chang
On Vision Features in Multimodal Machine Translation	Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma, JingBo Zhu
One Country, 700+ Languages: NLP Challenges for Underrepresented Languages and Dialects in Indonesia	Alham Fikri Aji, Genta Indra Winata, Fajri Koto, Samuel Cahyawijaya, Ade Romadhony, Rahmad Mahendra, Kemal Kurniawan, David Moeljadi, Radityo Eko Prasojo, Timothy Baldwin, Jey Han Lau, Sebastian Ruder
Online Semantic Parsing for Latency Reduction in Task-Oriented Dialogue	Jiawei Zhou, Jason Eisner, Michael Newman, Emmanouil Antonios Platanios, Sam Thomson
Open Domain Question Answering with A Unified Knowledge Interface	Kaixin Ma, Hao Cheng, Xiaodong Liu, Eric Nyberg, Jianfeng Gao
OpenHands: Making Sign Language Recognition Accessible with Pose-based Pretrained Models across Languages	Prem Selvaraj, Gokul NC, Pratyush Kumar, Mitesh M Khapra
Other Roles Matter! Enhancing Role-Oriented Dialogue Summarization via Role Interactions	Haitao Lin, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang, Chengqing Zong
Overcoming a Theoretical Limitation of Self-Attention	David Chiang, Peter Cholak
Overcoming Catastrophic Forgetting beyond Continual Learning: Balanced Training for Neural Machine Translation	Chenze Shao, Yang Feng
Overlap-based Vocabulary Generation Improves Cross-lingual Transfer Among Related Languages	Vaidehi Patil, Partha Talukdar, Sunita Sarawagi
Packed Levitated Marker for Entity and Relation Extraction	Deming Ye, Yankai Lin, Peng Li, Maosong Sun
ParaDetox: Detoxification with Parallel Data	Varvara Logacheva, Daryna Dementieva, Sergey Ustyantsev, Daniil Moskovskiy, David Dale, Irina Krotova, Nikita Semenov, Alexander Panchenko
Parallel Instance Query Network for Named Entity Recognition	Yongliang Shen, Xiaobin Wang, Zeqi Tan, Guangwei Xu, Pengjun Xie, Fei Huang, Weiming Lu, Yueting Zhuang
Pass off Fish Eyes for Pearls: Attacking Model Selection of Pre-trained Models	Biru Zhu, Yujia Qin, Fanchao Qi, Yangdong Deng, Zhiyuan Liu, Maosong Sun, Ming Gu
Perceiving the World: Question-guided Reinforcement Learning for Text-based Games	Yunqiu Xu, Meng Fang, Ling Chen, Yali Du, Joey Tianyi Zhou, Chengqi Zhang

Phone-ing it in: Towards Flexible Multi-Modal Language Model Training by Phonetic Representations of Data	Colin Leong, Daniel Lee Whitenack
Phrase-aware Unsupervised Constituency Parsing	Xiaotao Gu, Yikang Shen, Jiaming Shen, Jingbo Shang, Jiawei Han
PLANET: Dynamic Content Planning in Autoregressive Transformers for Long-form Text Generation	Zhe Hu, Hou Pong Chan, Jiachen Liu, Xinyan Xiao, Hua Wu, Lifu Huang
PPT: Pre-trained Prompt Tuning for Few-shot Learning	Yuxian Gu, Xu Han, Zhiyuan Liu, Minlie Huang
Pre-training and Fine-tuning Neural Topic Model: A Simple yet Effective Approach to Incorporating External Knowledge	Linhai Zhang, Xuemeng Hu, Boyu Wang, Deyu Zhou, Qian-Wen Zhang, Yunbo Cao
Pre-training to Match for Unified Low-shot Relation Extraction	Fangchao Liu, Hongyu Lin, Xianpei Han, Boxi Cao, Le Sun
Predicate-Argument Based Bi-Encoder for Paraphrase Identification	Qiwei Peng, David Weir, Julie Weeds, Yekun Chai
Predicting Intervention Approval in Clinical Trials through Multi-Document Summarization	Georgios Katsimpras, Georgios Paliouras
Prediction Difference Regularization against Perturbation for Neural Machine Translation	Dengji Guo, Zhengrui Ma, Min Zhang, Yang Feng
Premise-based Multimodal Reasoning: Conditional Inference on Joint Textual and Visual Clues	Qingxiu Dong, Ziwei Qin, Heming Xia, Tian Feng, Shoujie Tong, Haoran Meng, Lin Xu, zhongyu wei, Weidong Zhan, Baobao Chang, Sujian Li, Tianyu Liu, Zhifang Sui
Pretraining with Artificial Language: Studying Transferable Knowledge in Language Models	Ryokan Ri, Yoshimasa Tsuruoka
PRIMERA: Pyramid-based Masked Sentence Pre-training for Multi-document Summarization	Wen Xiao, Iz Beltagy, Giuseppe Carenini, Arman Cohan
Principled Paraphrase Generation with Parallel Corpora	Aitor Ormazabal, Mikel Artetxe, Aitor Soroa, Gorka Labaka, Eneko Agirre
Prix-LM: Pretraining for Multilingual Knowledge Base Construction	Wenxuan Zhou, Fangyu Liu, Ivan Vulić, Nigel Collier, Muhao Chen
Probing as Quantifying Inductive Bias	Alexander Immer, Lucas Torroba Hennigen, Vincent Fortuin, Ryan D Cotterell
Probing for Labeled Dependency Trees	Max Müller-Eberstein, Rob van der Goot, Barbara Plank
Probing for Predicate Argument Structures in Pretrained Language Models	Simone Conia, Roberto Navigli
Probing Simile Knowledge from Pre-trained Language Models	Weijie Chen, Yongzhu Chang, Rongsheng Zhang, Jiashu Pu, Guandan Chen, Le Zhang, Yadong Xi, Yijiang Chen, chang su
Probing Structured Pruning on Multilingual Pre-trained Models: Settings, Algorithms, and Efficiency	Yanyang Li, Fuli Luo, Runxin Xu, Songfang Huang, Fei Huang, Liwei Wang
Program Transfer for Answering Complex Questions over Knowledge Bases	Shulin Cao, Jiaxin Shi, Zijun Yao, Xin Lv, Jifan Yu, Lei Hou, Juanzi Li, Zhiyuan Liu, Jinghui Xiao

PromDA: Prompt-based Data Augmentation for Low-Resource NLU Tasks	Yufei Wang, Can Xu, Qingfeng Sun, Huang Hu, Chongyang Tao, Xiubo Geng, Daxin Jiang
Prompt for Extraction? PAIE: Prompting Argument Interaction for Event Argument Extraction	Yubo Ma, Zehao Wang, Yixin Cao, Mukai Li, Meiqi Chen, Kun Wang, Jing Shao
Prompt-Based Rule Discovery and Boosting for Interactive Weakly-Supervised Learning	Rongzhi Zhang, Yue Yu, Pranav Shetty, Le Song, Chao Zhang
Prompt-free and Efficient Few-shot Learning with Language Models	Rabeeh Karimi mahabadi, Luke Zettlemoyer, James Henderson, Lambert Mathias, Marzieh Saeidi, Veselin Stoyanov, Majid Yazdani
ProphetChat: Enhancing Dialogue Generation with Simulation of Future Conversation	Chang Liu, Xu Tan, Chongyang Tao, Zhenxin Fu, Dongyan Zhao, Tie-Yan Liu, Rui Yan
ProtoTEX: Explaining Model Decisions with Prototype Tensors	Anubrata Das, Chitrang Gupta, Venelin Kovatchev, Matthew Lease, Junyi Jessy Li
Prototypical Verbalizer for Prompt-based Few-shot Tuning	Ganqu Cui, Shengding Hu, Ning Ding, Longtao Huang, Zhiyuan Liu
Pyramid-BERT: Reducing Complexity via Successive Core-set based Token Selection	Xin Huang, Ashish Khetan, Rene Bidart, Zohar Karnin
QAConv: Question Answering on Informative Conversations	Chien-Sheng Wu, Andrea Madotto, Wenhao Liu, Pascale Fung, Caiming Xiong
Quality Controlled Paraphrase Generation	Elron Bandel, Ranit Aharonov, Michal Shmueli-Scheuer, Ilya Shnayderman, Noam Slonim, Liat Ein-Dor
Quantified Reproducibility Assessment of NLP Results	Anya Belz, Maja Popovic, Simon Mille
QuoteR: A Benchmark of Quote Recommendation for Writing	Fanchao Qi, Yanhui Yang, Jing Yi, Zhili Cheng, Zhiyuan Liu, Maosong Sun
Rare and Zero-shot Word Sense Disambiguation using Z-Reweighting	Ying Su, Hongming Zhang, Yangqiu Song, Tong Zhang
Rare Tokens Degenerate All Tokens: Improving Neural Text Generation via Adaptive Gradient Gating for Rare Token Embeddings	Sangwon Yu, Jongyoon Song, Heeseung Kim, Seongmin Lee, Woo-Jong Ryu, Sungroh Yoon
ReACC: A Retrieval-Augmented Code Completion Framework	Shuai Lu, Nan Duan, Hojae Han, Daya Guo, seung-won hwang, Alexey Svyatkovskiy
ReCLIP: A Strong Zero-Shot Baseline for Referring Expression Comprehension	Sanjay Subramanian, William Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, Anna Rohrbach
Redistributing Low-Frequency Words: Making the Most of Monolingual Data in Non-Autoregressive Translation	Liang Ding, Longyue Wang, Shuming Shi, Dacheng Tao, Zhaopeng Tu
Reducing Position Bias in Simultaneous Machine Translation with Length-Aware Framework	Shaolei Zhang, Yang Feng
Reinforcement Guided Multi-Task Learning Framework for Low-Resource Stereotype Detection	Rajkumar Pujari, Erik Oveson, Priyanka Kulkarni, Elnaz Nouri

RELIC: Retrieving Evidence for Literary Claims	Katherine Thai, Yapei Chang, Kalpesh Krishna, Mohit Iyyer
Reports of personal experiences and stories in argumentation: datasets and analysis	Neele Falk, Gabriella Lapesa
Requirements and Motivations of Low-Resource Speech Synthesis for Language Revitalization	Aidan Pine, Dan Wells, Nathan Brinklow, Patrick William Littell, Korin Richmond
Rethinking Negative Sampling for Handling Missing Entity Annotations	Yangming Li, Lemao Liu, Shuming Shi
Rethinking Self-Supervision Objectives for Generalizable Coherence Modeling	Prathyusha Jwalapuram, Shafiq Joty, Xiang Lin
Retrieval-guided Counterfactual Generation for QA	Bhargavi Paranjape, Matthew Lamm, Ian Tenney
Revisiting Over-Smoothness in Text to Speech	Yi Ren, Xu Tan, Tao Qin, Zhou Zhao, Tie-Yan Liu
Rewire-then-Probe: A Contrastive Recipe for Probing Biomedical Knowledge of Pre-trained Language Models	Zaiqiao Meng, Fangyu Liu, Ehsan Shareghi, Yixuan Su, Charlotte Anne Collins, Nigel Collier
Right for the Right Reason: Evidence Extraction for Trustworthy Tabular Reasoning	Vivek Gupta, Shuo Zhang, Alakananda Vempala, Yujie He, Temma Choji, Vivek Srikumar
RNG-KBQA: Generation Augmented Iterative Ranking for Knowledge Base Question Answering	Xi Ye, Semih Yavuz, Kazuma Hashimoto, Yingbo Zhou, Caiming Xiong
RNSum: A Large-Scale Dataset for Automatic Release Note Generation via Commit Logs Summarization	Hisashi Kamezawa, Noriki Nishida, Nobuyuki Shimizu, Takashi Miyazaki, Hideki Nakayama
Robust Lottery Tickets for Pre-trained Language Models	Rui Zheng, Bao Rong, Yuhao Zhou, Di Liang, Sirui Wang, Wei Wu, Tao Gui, Qi Zhang, Xuanjing Huang
RoCBert: Robust Chinese Bert with Multimodal Contrastive Pretraining	Hui Su, Weiwei Shi, Xiaoyu Shen, Zhou Xiao, Tuo Ji, Jiarui Fang, Jie Zhou
RoMe: A Robust Metric for Evaluating Natural Language Generation	Md Rashad Al Hasan Rony, Liubov Kovriguina, Debanjan Chaudhuri, Ricardo Usbeck, Jens Lehmann
RotateQVS: Representing Temporal Information as Rotations in Quaternion Vector Space for Temporal Knowledge Graph Completion	Kai Chen, Ye Wang, Yitong Li, Aiping Li
RST Discourse Parsing with Second-Stage EDU-Level Pre-training	Nan Yu, Meishan Zhang, Guohong Fu, Min Zhang
SaFeRDialogues: Taking Feedback Gracefully after Conversational Safety Failures	Megan Ung, Jing Xu, Y-Lan Boureau
SafetyKit: First Aid for Measuring Safety in Open-domain Conversational Systems	Emily Dinan, Gavin Abercrombie, A. Stevie Bergman, Shannon L. Spruit, Dirk Hovy, Y-Lan Boureau, Verena Rieser
SalesBot: Transitioning from Chit-Chat to Task-Oriented Dialogues	Ssu Chiu, Maolin Li, Yen-Ting Lin, Yun-Nung Chen
Saliency as Evidence: Event Detection with Trigger Saliency Attribution	Jian Liu, Yufeng Chen, Jinan Xu
Scheduled Multi-task Learning for Neural Chat Translation	Yunlong Liang, Fandong Meng, Jinan Xu, Yufeng

	Chen, Jie Zhou
SciNLI: A Corpus for Natural Language Inference on Scientific Text	Mobashir Sadat, Cornelia Caragea
SDR: Efficient Neural Re-ranking using Succinct Document Representation	Nachshon Cohen, Amit Portnoy, Besnik Fetahu, Amir Ingber
Searching for fingerspelled content in American Sign Language	Bowen Shi, Diane Brentari, Greg Shakhnarovich, Karen Livescu
Self-supervised Semantic-driven Phoneme Discovery for Zero-resource Speech Recognition	Liming Wang, Siyuan Feng, Mark A. Hasegawa-Johnson, Chang D. Yoo
Semantic Composition with PSHRG for Derivation Tree Reconstruction from Graph-Based Meaning Representations	Chun Hei Lo, Wai Lam, Hong Cheng
Semi-supervised Domain Adaptation for Dependency Parsing with Dynamic Matching Network	Ying Li, Shuaike Li, Min Zhang
Semi-Supervised Formality Style Transfer with Consistency Training	Ao Liu, An Wang, Naoaki Okazaki
Sense Embeddings are also Biased -- Evaluating Social Biases in Static and Contextualised Sense Embeddings	Yi Zhou, Masahiro Kaneko, Danushka Bollegala
Sentence-aware Contrastive Learning for Open-Domain Passage Retrieval	Wu Bo Hong, Zhuosheng Zhang, Jinyuan Wang, hai zhao
Sentence-level Privacy for Document Embeddings	Casey Meehan, Khalil Mrini, Kamalika Chaudhuri
Sequence-to-Sequence Knowledge Graph Completion and Question Answering	Apoorv Umang Saxena, Adrian Kochsiek, Rainer Gemulla
Sharpness-Aware Minimization Improves Language Model Generalization	Dara Bahri, Hossein Mobahi, Yi Tay
SHIELD: Defending Textual Neural Networks against Multiple Black-Box Adversarial Attacks with Stochastic Multi-Expert Patcher	Thai Le, Noseong Park, Dongwon Lee
Should a Chatbot be Sarcastic? Understanding User Preferences Towards Sarcasm Generation	Silviu Vlad Oprea, Steven R. Wilson, Walid Magdy
Show Me More Details: Discovering Hierarchies of Procedures from Semi-structured Web Data	Shuyan Zhou, Li Zhang, Yue Yang, Qing Lyu, Pengcheng Yin, Chris Callison-Burch, Graham Neubig
Signal in Noise: Exploring Meaning Encoded in Random Character Sequences with Character-Aware Language Models	Mark Bo Chu, Bhargav Srinivasa Desikan, Ethan O Nadler, Donald Ruggiero Lo Sardo, Elise Darragh-Ford, Douglas Richard Guilbeault
SimKGC: Simple Contrastive Knowledge Graph Completion with Pre-trained Language Models	Liang Wang, wei zhao, Zhuoyu Wei, Jingming Liu
Simulating Bandit Learning from User Feedback for Extractive Question Answering	Ge Gao, Eunsol Choi, Yoav Artzi
Situated Dialogue Learning through Procedural Environment Generation	Prithviraj Ammanabrolu, Renee Jia, Mark Riedl
Skill Induction and Planning with Latent Language	Pratyusha Sharma, Antonio Torralba, Jacob Andreas

SkipBERT: Efficient Inference with Shallow Layer Skipping	Jue WANG, Ke Chen, Gang Chen, Lidan Shou, Julian McAuley
Slangvolution: A Causal Analysis of Semantic Change and Frequency Dynamics in Slang	Daphna Keidar, Andreas Opedal, Zhijing Jin, Mrinmaya Sachan
So Different Yet So Alike! Constrained Unsupervised Text Style Transfer	Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan, Roger Zimmermann, Soujanya Poria
Softmax Bottleneck Makes Language Models Unable to Represent Multi-mode Word Distributions	Haw-Shiuan Chang, Andrew McCallum
Sparse Progressive Distillation: Resolving Overfitting under Pretrain-and-Finetune Paradigm	Shaoyi Huang, Dongkuan Xu, Ian En-Hsu Yen, Yijue Wang, Sung-En Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Sanguthevar Rajasekaran, Hang Liu, Caiwen Ding
Sparsifying Transformer Models with Trainable Representation Pooling	Michał Pietruszka, Łukasz Borchmann, Łukasz Garncarek
Speaker Information Can Guide Models to Better Inductive Biases: A Case Study On Predicting Code-Switching	Alissa Ostapenko, Shuly Wintner, Melinda Fricke, Yulia Tsvetkov
SpeechT5: Unified-Modal Encoder-Decoder Pre-Training for Spoken Language Processing	Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie LIU, Tom Ko, Qing Li, Yu Zhang, Zhihua Wei, Yao Qian, Jinyu Li, Furu Wei
SPoT: Better Frozen Model Adaptation through Soft Prompt Transfer	Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou', Daniel Cer
Spurious Correlations in Reference-Free Evaluation of Text Generation	Esin Durmus, Faisal Ladhak, Tatsunori Hashimoto
SRL4E – Semantic Role Labeling for Emotions: A Unified Evaluation Framework	Cesare Campagnano, Simone Conia, Roberto Navigli
StableMoE: Stable Routing Strategy for Mixture of Experts	Damai Dai, Li Dong, Shuming Ma, Bo Zheng, Zhifang Sui, Baobao Chang, Furu Wei
STEMM: Self-learning with Speech-text Manifold Mixup for Speech Translation	Qingkai Fang, Rong Ye, Lei Li, Yang Feng, Mingxuan Wang
Structural Characterization for Dialogue Disentanglement	Xinbei Ma, Zhuosheng Zhang, hai zhao
Structured Pruning Learns Compact and Accurate Models	Mengzhou Xia, Zexuan Zhong, Danqi Chen
Subgraph Retrieval Enhanced Model for Multi-hop Knowledge Base Question Answering	Jing Zhang, Xiaokang Zhang, Jifan Yu, Jian Tang, Jie Tang, Cuiping Li, Hong Chen
Substructure Distribution Projection for Zero-Shot Cross-Lingual Dependency Parsing	Freda Shi, Kevin Gimpel, Karen Livescu
Summ^N: A Multi-Stage Summarization Framework for Long Input Dialogues and Documents	Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Hassan Awadallah, Dragomir Radev, Rui Zhang
SummaReranker: A Multi-Task Mixture-of-Experts Re-ranking Framework	Mathieu Ravaut, Shafiq Joty, Nancy F. Chen

for Abstractive Summarization	
SummScreen: A Dataset for Abstractive Screenplay Summarization	Mingda Chen, Zewei Chu, Sam Wiseman, Kevin Gimpel
SUPERB-SG: Enhanced Speech processing Universal PERFORMANCE Benchmark for Semantic and Generative Capabilities	Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhota, Shu-wen Yang, Shuyan Dong, Andy T. Liu, Cheng-I Lai, Jiatong Shi, Xuankai Chang, Phil Hall, Hsuan-Jui Chen, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, Hung-yi Lee
Synthetic Question Value Estimation for Domain Adaptation of Question Answering	Xiang Yue, Ziyu Yao, Huan Sun
Systematic Inequalities in Language Technology Performance across the World's Languages	Damian E Blasi, Antonios Anastasopoulos, Graham Neubig
TableFormer: Robust Transformer Modeling for Table-Text Encoding	Jingfeng Yang, Aditya Gupta, Shyam Upadhyay, Luheng He, Rahul Goel, Shachi Paul
Tackling Fake News Detection by Continually Improving Social Context Representations using Graph Neural Networks	Nikhil Mehta, Maria Leonor Pacheco, Dan Goldwasser
Tailor: Generating and Perturbing Text with Semantic Controls	Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E Peters, Matt Gardner
Text-Free Prosody-Aware Generative Spoken Language Modeling	Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhota, Tu-Anh Nguyen, Morgane Rivière, Abdelrahman Mohamed, Emmanuel Dupoux, Wei-Ning Hsu
Text-to-Table: A New Way of Information Extraction	Xueqing Wu, Jiacheng Zhang, Hang Li
Textomics: A Dataset for Genomics Data Summary Generation	Mu-Chun Wang, Zixuan Liu, Sheng Wang
That Is a Suspicious Reaction!: Interpreting Logits Variation to Detect NLP Adversarial Attacks	Edoardo Mosca, Shreyash Agarwal, Javier Rando Ramírez, Georg Groh
That Slep en AI the Nyght with Open Ye! Cross-era Sequence Segmentation with Switch-memory	Xuemei Tang, Qi Su
The AI Doctor Is In: A Survey of Task-Oriented Dialogue Systems for Healthcare Applications	Mina Valizadeh, Natalie Parde
The Dangers of Underclaiming: Reasons for Caution When Reporting How NLP Systems Fail	Samuel R. Bowman
The Grammar-Learning Trajectories of Neural Language Models	Leshem Choshen, Guy Hachohen, Daphna Weinshall, Omri Abend
The Moral Debater: A Study on the Computational Generation of Morally Framed Arguments	Milad Alshomary, Roxanne El Baff, Timon Gurcke, Henning Wachsmuth
The Moral Integrity Corpus: A Benchmark for Ethical Dialogue Systems	Caleb Ziems, Jane A. Yu, Yi-Chia Wang, Alon Y. Halevy, Diyi Yang
The Paradox of the Compositionality of Natural Language: A Neural Machine Translation Case Study	Verna Dankers, Elia Bruni, Dieuwke Hupkes

The patient is more dead than alive: exploring the current state of the multi-document summarisation of the biomedical literature	Yulia Otmakhova, Karin Verspoor, Timothy Baldwin, Jey Han Lau
The Trade-offs of Domain Adaptation for Neural Language Models	David Grangier, Dan Iter
There Are a Thousand Hamlets in a Thousand People's Eyes: Enhancing Knowledge-grounded Dialogue with Personal Memory	Tingchen Fu, Xueliang Zhao, Chongyang Tao, Ji-Rong Wen, Rui Yan
There's a Time and Place for Reasoning Beyond the Image	Xingyu Fu, Ben Zhou, Ishaan Preetam Chandratreya, Carl Vondrick, Dan Roth
Things not Written in Text: Exploring Spatial Commonsense from Visual Signals	Xiao Liu, Da Yin, Yansong Feng, Dongyan Zhao
Think Before You Speak: Explicitly Generating Implicit Commonsense Knowledge for Response Generation	Pei Zhou, Karthik Gopalakrishnan, Behnam Hedayatnia, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu, Dilek Hakkani-Tur
Token Dropping for Efficient BERT Pretraining	Le Hou, Richard Yuanzhe Pang, Tianyi Zhou, Yuexin Wu, Xinying Song, Xiaodan Song, Denny Zhou
TopWORDS-Seg: Simultaneous Text Segmentation and Word Discovery for Open-Domain Chinese Texts via Bayesian Inference	Changzai Pan, Maosong Sun, Ke Deng
Toward Annotator Group Bias in Crowdsourcing	Haochen Liu, Joseph Thekinen, Sinem Mollaoglu, Da Tang, Ji Yang, Youlong Cheng, Hui Hui, Jiliang Tang
Toward Interpretable Semantic Textual Similarity via Optimal Transport-based Contrastive Sentence Learning	Seonghyeon Lee, Dongha Lee, Seongbo Jang, Hwanjo Yu
Towards Abstractive Grounded Summarization of Podcast Transcripts	Kaiqiang Song, Chen Li, Xiaoyang Wang, Dong Yu, Fei Liu
Towards Afrocentric NLP for African Languages: Where We Are and Where We Can Go	Ife Adebara, Muhammad Abdul-Mageed
Towards Better Characterization of Paraphrases	Timothy Liu, De Wen Soh
Towards Comprehensive Patent Approval Predictions: Beyond Traditional Document Classification	Xiaochen Kev Gao, Zhaoyi Hou, Yifei Ning, Kewen Zhao, Beilei He, Jingbo Shang, Vish Krishnan
Towards Learning (Dis-)Similarity of Source Code from Program Contrasts	Yangruibo Ding, Luca Buratti, Saurabh Pujar, Alessandro Morari, Baishakhi Ray, Saikat Chakraborty
Towards Making the Most of Cross-Lingual Transfer for Zero-Shot Neural Machine Translation	Guanhua Chen, Shuming Ma, Yun Chen, Dongdong Zhang, Jia Pan, Wenping Wang, Furu Wei
Towards Robustness of Text-to-SQL Models Against Natural and Realistic Adversarial Table Perturbation	Xinyu Pi, Bing Wang, Yan Gao, Jiaqi Guo, Zhoujun Li, Jian-Guang Lou
ToxiGen: A Large-Scale Machine-Generated Dataset for Adversarial and Implicit Hate Speech Detection	Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, Ece Kamar
Training Data is More Valuable than You Think: A Simple and Effective	Shuohang Wang, Yichong Xu, Yuwei Fang, Yang

Method by Retrieving from Training Data	Liu, Siqi Sun, Ruochen Xu, Chenguang Zhu, Michael Zeng
Transformers in the loop: Polarity in neural models of language	Lisa Bylina, Alexey Tikhonov
Transkimmer: Transformer Learns to Layer-wise Skim	Yue Guan, Zhengyi Li, Jingwen Leng, Zhouhan Lin, Minyi Guo
TruthfulQA: Measuring How Models Mimic Human Falsehoods	Stephanie Lin, Jacob Hilton, Owain Evans
Turning Tables: Generating Examples from Semi-structured Tables for Endowing Language Models with Reasoning Skills	Ori Yoran, Alon Talmor, Jonathan Berant
Twittlrish: A Universal Dependencies Treebank of Tweets in Modern Irish	Lauren Cassidy, Teresa Lynn, James Barry, Jennifer Foster
UCTopic: Unsupervised Contrastive Learning for Phrase Representations and Topic Mining	Jiacheng Li, Jingbo Shang, Julian McAuley
Uncertainty Determines the Adequacy of the Mode and the Tractability of Decoding in Sequence-to-Sequence Models	Felix Stahlberg, Ilia Kulikov, Shankar Kumar
Uncertainty Estimation of Transformer Predictions for Misclassification Detection	Artem Vazhentsev, Gleb Kuzmin, Artem Shelmanov, Akim Tsvigun, Evgenii Tsymbalov, Kirill Fedyanin, Maxim Panov, Alexander Panchenko, Gleb Gennadjevich Gusev, Mikhail Burtsev, Manvel Avetisian, Leonid Zhukov
Under the Morphosyntactic Lens: A Multifaceted Evaluation of Gender Bias in Speech Translation	Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, Marco Turchi
Understanding and Improving Sequence-to-Sequence Pretraining for Neural Machine Translation	Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu, Michael Lyu
Understanding Gender Bias in Knowledge Base Embeddings	Yupei Du, Qi Zheng, Yuanbin Wu, Man Lan, Yan Yang, Meirong Ma
Understanding Iterative Revision from Human-Written Text	Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, Dongyeop Kang
Understanding Multimodal Procedural Knowledge by Sequencing Multimodal Instructional Manuals	Te-Lin Wu, Alex Spangher, Pegah Alipoormolabashi, Marjorie Freedman, Ralph M. Weischedel, Nanyun Peng
Unified Speech-Text Pre-training for Speech Translation and Recognition	Yun Tang, Hongyu Gong, Ning Dong, Changhan Wang, Wei-Ning Hsu, Jiatao Gu, Alexei Baevski, Xian Li, Abdelrahman Mohamed, Michael Auli, Juan Pino
Unified Structure Generation for Universal Information Extraction	Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, Hua Wu
UniPELT: A Unified Framework for Parameter-Efficient Language Model Tuning	Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Scott Yih, Madian Khabsa
UniTE: Unified Translation Evaluation	Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek F. Wong, Lidia S.

	Chao
UniTranSeR: A Unified Transformer Semantic Representation Framework for Multimodal Task-Oriented Dialog System	Zhiyuan Ma, Jianjun Li, Guohui Li, Yongjing Cheng
Universal Conditional Masked Language Pre-training for Neural Machine Translation	Pengfei Li, Liangyou Li, Meng Zhang, Minghao Wu, Qun Liu
UniXcoder: Unified Cross-Modal Pre-training for Code Representation	Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, Jian Yin
Unsupervised Corpus Aware Language Model Pre-training for Dense Passage Retrieval	Luyu Gao, Jamie Callan
Unsupervised Dependency Graph Network	Yikang Shen, Shawn Tan, Alessandro Sordoni, Peng Li, Jie Zhou, Aaron Courville
Unsupervised Extractive Opinion Summarization Using Sparse Coding	Somnath Basu Roy Chowdhury, Chao Zhao, Snigdha Chaturvedi
Updated Headline Generation: Creating Updated Summaries for Evolving News Stories	Sheena Panthaplackel, Adrian Benton, Mark Dredze
Upstream Mitigation Is Not All You Need: Testing the Bias Transfer Hypothesis in Pre-Trained Language Models	Ryan Steed, Swetasudha Panda, Ari Kobren, Michael L Wick
Using Context-to-Vector with Graph Retrofitting to Improve Word Embeddings	Jiangbin Zheng, Yile Wang, Ge Wang, Jun Xia, Yufei Huang, Guojiang Zhao, Yue Zhang, Stan Z. Li
VALSE: A Task-Independent Benchmark for Vision and Language Models Centered on Linguistic Phenomena	Letitia Parcalabescu, Michele Cafagna, Lilita Muradjan, Anette Frank, Iacer Calixto, Albert Gatt
VALUE: Understanding Dialect Disparity in NLU	Caleb Ziems, Jiaao Chen, Camille Harris, Jessica Anderson, Diyi Yang
Variational Graph Autoencoding as Cheap Supervision for AMR Coreference Resolution	Irene Li, Linfeng Song, Kun Xu, Dong Yu
Vision-and-Language Navigation: A Survey of Tasks, Methods, and Future Directions	Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason, Xin Eric Wang
Vision-Language Pre-Training for Multimodal Aspect-Based Sentiment Analysis	Yan Ling, Jianfei Yu, Rui Xia
Visual-Language Navigation Pretraining via Prompt-based Environmental Self-exploration	Xiwen Liang, Fengda Zhu, Li Lingling, Hang Xu, Xiaodan Liang
WatClaimCheck: A new Dataset for Claim Entailment and Inference	Kashif Khan, Ruizhe Wang, Pascal Poupart
Weakly Supervised Word Segmentation for Computational Language Documentation	Shu Okabe, Laurent BesacierUGA, François Yvon
What does the sea say to the shore? A BERT based DST style approach for speaker to dialogue attribution in novels	Carolina Cuesta-Lazaro, Animesh Prasad, Trevor Wood
What Makes Reading Comprehension Questions Difficult?	Saku Sugawara, Nikita Nangia, Alex Warstadt, Samuel R. Bowman

When did you become so smart, oh wise one?! Sarcasm Explanation in Multi-modal Multi-party Dialogues	Shivani Kumar, Atharva Kulkarni, Md Shad Akhtar, Tanmoy Chakraborty
Where to Go for the Holidays: Towards Mixed-Type Dialogs for Clarification of User Goals	Zeming Liu, Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu
Which side are you on? Insider-Outsider classification in conspiracy-theoretic social media	Pavan S Holur, Tianyi Wang, Shadi Shahsavari, Timothy R Tangherlini, Vwani P Roychowdhury
WikiDiverse: A Multimodal Entity Linking Dataset with Diversified Contextual Topics and Entity Types	Xuwu Wang, Junfeng Tian, Min Gui, Zhixu Li, Rui Wang, Ming Yan, Lihan Chen, Yanghua Xiao
Word Order Does Matter and Shuffled Language Models Know It	Mostafa Abdou, Vinit Ravishankar, Artur Kulmizev, Anders Søgaard
Word Segmentation as Unsupervised Constituency Parsing	Raquel G. Alhama
Word2Box: Capturing Set-Theoretic Semantics of Words using Box Embeddings	Shib Sankar Dasgupta, Michael Boratko, Siddhartha Mishra, Shriya Atmakuri, Dhruvesh Patel, Xiang Lorraine Li, Andrew McCallum
XLM-E: Cross-lingual Language Model Pre-training via ELECTRA	Zewen Chi, Shaohan Huang, Li Dong, Shuming Ma, Bo Zheng, Saksham Singhal, Payal Bajaj, Xia Song, Xian-Ling Mao, Heyan Huang, Furu Wei
You might think about slightly revising the title: Identifying Hedges in Peer-tutoring Interactions	Yann Raphalen, Chloé Clavel, Justine Cassell
Your Answer is Incorrect... Would you like to know why? Introducing a Bilingual Short Answer Feedback Dataset	Anna Filighera, Siddharth Singh Parihar, Tim Steuer, Tobias Meuser, Sebastian Ochs
Zero-Shot Cross-lingual Semantic Parsing	Tom Sherborne, Mirella Lapata
Zoom Out and Observe: News Environment Perception for Fake News Detection	Qiang Sheng, Juan Cao, Xueyao Zhang, Rundong Li, Danding Wang, Yongchun Zhu

Short Papers

Title	Authors
A Copy-Augmented Generative Model for Open-Domain Question Answering	Shuang Liu, Dong Wang, Xiaoguang Li, Minghui Huang, Meizhen Ding
A Flexible Multi-Task Model for BERT Serving	Tianwen Wei, Jianwei Qi, Shenghuan He
A Recipe for Arbitrary Text Style Transfer with Large Language Models	Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, Jason Wei
A Risk-Averse Mechanism for Suicidality Assessment on Social Media	Ramit Sawhney, Atula Tejaswi Neerkaje, Manas Gaur
A Simple but Effective Pluggable Entity Lookup Table for Pre-trained Language Models	Deming Ye, Yankai Lin, Peng Li, Maosong Sun, Zhiyuan Liu
Adjusting the Precision-Recall Trade-Off with Align-and-Predict Decoding for Grammatical Error Correction	Xin Sun, Houfeng Wang

An Analysis of Negation in Natural Language Understanding Corpora	Md Mosharaf Hossain, Dhivya Chinnappa, Eduardo Blanco
An Embarrassingly Simple Method to Mitigate Undesirable Properties of Pretrained Language Model Tokenizers	Valentin Hofmann, Hinrich Schuetze, Janet B. Pierrehumbert
Analyzing Wrap-Up Effects through an Information-Theoretic Lens	Clara Isabel Meister, Tiago Pimentel, Thomas Hikaru Clark, Ryan D Cotterell, Roger P. Levy
Are Shortest Rationales the Best Explanations for Human Understanding?	Hua Shen, Tongshuang Wu, Wenbo Guo, Ting-Hao Huang
As Little as Possible, as Much as Necessary: Detecting Over- and Undertranslations with Contrastive Conditioning	Jannis Vamvas, Rico Sennrich
Augmenting Document Representations for Dense Retrieval with Interpolation and Perturbation	Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, Jong C. Park
Automatic Detection of Entity-Manipulated Text using Factual Knowledge	Ganesh Jawahar, Muhammad Abdul-Mageed, Laks V. S. Lakshmanan
BitFit: Simple Parameter-efficient Fine-tuning for Transformer-based Masked Language-models	Elad Ben Zaken, Yoav Goldberg, Shauli Ravfogel
Buy Tesla, Sell Ford: Assessing Implicit Stock Market Preference in Pre-trained Language Models	Chengyu Chuang, Yi Yang
C-MORE: Pretraining to Answer Open-Domain Questions by Consulting Millions of References	Xiang Yue, Xiaoman Pan, Wenlin Yao, Dian Yu, Dong Yu, Jianshu Chen
Can a Transformer Pass the Wug Test? Tuning Copying Bias in Neural Morphological Inflection Models	Ling Liu, Mans Hulden
Can Visual Dialogue Models Do Scorekeeping? Exploring How Dialogue Representations Incrementally Encode Shared Knowledge	Brielen Madureira, David Schlangen
Canary Extraction in Natural Language Understanding Models	Rahil Parikh, Christophe Dupuy, Rahul Gupta
CoDA21: Evaluating Language Understanding Capabilities of NLP Models With Context-Definition Alignment	Lütfi Kerem Senel, Timo Schick, Hinrich Schuetze
Code Synonyms Do Matter: Multiple Synonyms Matching Network for Automatic ICD Coding	Zheng Yuan, Chuanqi Tan, Songfang Huang
Complex Evolutional Pattern Learning for Temporal Knowledge Graph Reasoning	Zixuan Li, Saiping Guan, Xiaolong Jin, Weihua Peng, Yajuan Lyu, Yong Zhu, Long Bai, Wei Li, Jiafeng Guo, Xueqi Cheng
Contrastive Learning-Enhanced Nearest Neighbor Mechanism for Multi-Label Text Classification	Xi'ao Su, Ran Wang, Xinyu Dai
Counterfactual Explanations for Natural Language Interfaces	George Tolkachev, Stephen Mell, Stephan Zdanczewic, Osbert Bastani
Data Contamination: From Memorization to Exploitation	Inbal Magar, Roy Schwartz
Detecting Annotation Errors in Morphological Data with the Transformer	Ling Liu, Mans Hulden
Developmental Negation Processing in Transformer Language Models	Antonio Laverghetta Jr., John Licato

Direct parsing to sentiment graphs	David Samuel, Jeremy Barnes, Robin Kurtz, Stephan Oepen, Lilja Øvrelid, Erik Velldal
DiS-ReX: A Multilingual Dataset for Distantly Supervised Relation Extraction	Abhyuday Bhartiya, Kartikeya Badola, Mausam .
Disentangled Knowledge Transfer for OOD Intent Discovery with Unified Contrastive Learning	Yutao Mou, Keqing He, Yanan Wu, Zhiyuan Zeng, Hong Xu, Huixing Jiang, Wei Wu, Weiran Xu
DMix: Adaptive Distance-aware Interpolative Mixup	Ramit Sawhney, Megh Thakkar, Shrey Pandit, Ritesh Singh Soun, Di Jin, Diyi Yang, Lucie Flek
Does BERT Know that the IS-A Relation Is Transitive?	Ruixi Lin, Hwee Tou Ng
DQ-BART: Efficient Sequence-to-Sequence Model via Joint Distillation and Quantization	Zheng Li, Zijian Wang, Ming Tan, Ramesh Nallapati, Parminder Bhatia, Andrew Arnold, Bing Xiang, Dan Roth
Efficient Classification of Long Documents Using Transformers	Hyunji Hayley Park, Yogarshi Vyas, Kashif Shah
Estimating the Entropy of Linguistic Distributions	Aryaman Arora, Clara Isabel Meister, Ryan D Cotterell
Event-Event Relation Extraction using Probabilistic Box Embedding	EunJeong Hwang, Jay-Yoon Lee, Tianyi Yang, Dhruvesh Patel, Dongxu Zhang, Andrew McCallum
Exploiting Language Model Prompts Using Similarity Measures: A Case Study on the Word-in-Context Task	Mohsen Tabasi, Kiamehr Rezaee, Mohammad Taher Pilehvar
Fire Burns, Sword Cuts: Commonsense Inductive Bias for Exploration in Text-based Games	Dongwon Kelvin Ryu, Ehsan Shareghi, Meng Fang, Yunqiu Xu, Shirui Pan, Reza Haf
Focus on the Target's Vocabulary: Masked Label Smoothing for Machine Translation	Liang Chen, Runxin Xu, Baobao Chang
Have my arguments been replied to? Argument Pair Extraction as Machine Reading Comprehension	Jianzhu Bao, Jingyi Sun, Qinglin Zhu, Ruifeng Xu
Hierarchical Curriculum Learning for AMR Parsing	Peiyi Wang, Liang Chen, Tianyu Liu, Damai Dai, Yunbo Cao, Baobao Chang, Zhifang Sui
High probability or low information? The probability–quality paradox in language generation	Clara Isabel Meister, Gian Wiher, Tiago Pimentel, Ryan D Cotterell
How Distributed are Distributed Representations? An Observation on the Locality of Syntactic Information in Verb Agreement Tasks	Bingzhi Li, Guillaume Wisniewski, Benoit Crabbé
How does the pre-training objective affect what large language models learn about linguistic properties?	Ahmed Alajrami, Nikolaos Aletras
How reparametrization trick broke differentially-private text representation learning	Ivan Habernal
HYPHEN: Hyperbolic Hawkes Attention For Text Streams	Shivam Agarwal, Ramit Sawhney, Sanchit Ahuja, Ritesh Singh Soun, Sudheer Chava
Investigating person-specific errors in chat-oriented dialogue systems	Koh Mitsuda, Ryuichiro Higashinaka, Tingxuan Li, Sen Yoshida

k-Rater Reliability: The Correct Unit of Reliability for Aggregated Human Annotations	Ka Wong, Praveen Paritosh
Kronecker Decomposition for GPT Compression	Ali Edalati, Marzieh S. Tahaei, Ahmad Rashid, Vahid Partovi Nia, James J. Clark, Mehdi Rezagholizadeh
Learning-by-Narrating: Narrative Pre-Training for Zero-Shot Dialogue Comprehension	Chao Zhao, Wenlin Yao, Dian Yu, Kaiqiang Song, Dong Yu, Jianshu Chen
Leveraging Explicit Lexico-logical Alignments in Text-to-SQL Parsing	Runxin Sun, Shizhu He, Chong Zhu, Yaohan He, Jinlong Li, Jun Zhao, Kang Liu
LM-BFF-MS: Improving Few-Shot Fine-tuning of Language Models based on Multiple Soft Demonstration Memory	Eunhwan Park, Donghyeon Jeon, Seonhoon Kim, Inho Kang, Seung-Hoon Na
Machine Translation for Livonian: Catering to 20 Speakers	Matīss Rikters, Marili Tomingas, Tuuli Tuisk, Valts Ernštreits, Mark Fishel
Mismatch between Multi-turn Dialogue and its Evaluation Metric in Dialogue State Tracking	Takyong Kim, Hoonsang Yoon, Yukyung Lee, Pilsung Kang, Misuk Kim
Morphological Reinflection with Multiple Arguments: An Extended Annotation schema and a Georgian Case Study	David Guriel, Omer Goldman, Reut Tsarfaty
Multilingual Pre-training with Language and Task Adaptation for Multilingual Text Style Transfer	Huiyuan Lai, Antonio Toral, Malvina Nissim
NoisyTune: A Little Noise Can Help You Finetune Pretrained Language Models Better	Chuhan Wu, Fangzhao Wu, Tao Qi, Yongfeng Huang
On Efficiently Acquiring Annotations for Multilingual Models	Joel Ruben Antony Moniz, Barun Patra, Matthew R. Gormley
On the Effect of Isotropy on VAE Representations of Text	Lan Zhang, Wray Buntine, Ehsan Shareghi
On the Importance of Effectively Adapting Pretrained Language Models for Active Learning	Katerina Margatina, Loic Barrault, Nikolaos Aletras
On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations	Yang Trista Cao, Yada Pruksachatkun, Kai-Wei Chang, Rahul Gupta, Varun Kumar, Jwala Dhamala, Aram Galstyan
P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks	Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, Jie Tang
PARE: A Simple and Strong Baseline for Monolingual and Multilingual Distantly Supervised Relation Extraction	Vipul Kumar Rathore, Kartikeya Badola, Parag Singla, Mausam .
Pixie: Preference in Implicit and Explicit Comparisons	Amanul Haque, Vaibhav Garg, Hui Guo, Munindar P. Singh
Predicting Difficulty and Discrimination of Natural Language Questions	Matthew Alexander Byrd, Shashank Srivastava
Predicting Sentence Deletions for Text Simplification Using a Functional Discourse Structure	Bohan Zhang, Prafulla Kumar Choubey, Ruihong Huang
PriMock57: A Dataset Of Primary Care Mock Consultations	Alex Papadopoulos Korfiatis, Francesco Moramarco, Radmila Sarac, Aleksandar Savkov

Primum Non Nocere: Before working with Indigenous data, the ACL must confront ongoing colonialism	Lane Schwartz
Probing the Robustness of Trained Metrics for Conversational Dialogue Systems	Jan Milan Deriu, Don Tuggener, Pius von Däniken, Mark Cieliebak
Problems with Cosine as a Measure of Embedding Similarity for High Frequency Words	Kaitlyn Zhou, Kawin Ethayarajh, Dallas Card, Dan Jurafsky
Region-dependent temperature scaling for certainty calibration and application to class-imbalanced token classification	Hillary Dawkins, Isar Nejadgholi
Rethinking and Refining the Distinct Metric	Siyang Liu, Sahand Sabour, Yinhe Zheng, Pei Ke, Xiaoyan Zhu, Minlie Huang
Revisiting the Compositional Generalization Abilities of Neural Sequence Models	Arkil Patel, Satwik Bhattamishra, Phil Blunsom, Navin Goyal
Rewarding Semantic Similarity under Optimized Alignments for AMR-to-Text Generation	Lisa Jin, Daniel Gildea
S ⁴ -Tuning: A Simple Cross-lingual Sub-network Tuning Method	Runxin Xu, Fuli Luo, Baobao Chang, Songfang Huang, Fei Huang
Sample, Translate, Recombine: Leveraging Audio Alignments for Data Augmentation in End-to-end Speech Translation	Tsz Kin Lam, Shigehiko Schamoni, Stefan Riezler
SCD: Self-Contrastive Decorrelation of Sentence Embeddings	Tassilo Klein, Moin Nabi
Sequence-to-sequence AMR Parsing with Ancestor Information	Chen Yu, Daniel Gildea
Simple and Effective Knowledge-Driven Query Expansion for QA-Based Product Attribute Extraction	Keiji Shinzato, Naoki Yoshinaga, Yandi Xia, Wei-Te Chen
Sub-Word Alignment is Still Useful: A Vest-Pocket Method for Enhancing Low-Resource Machine Translation	Minhan Xu, Yu Hong
Text Smoothing: Enhance Various Data Augmentation Methods on Text Classification Tasks	Xing W, Chaochen Gao, Meng Lin, Liangjun Zang, Songlin Hu
The Power of Prompt Tuning for Low-Resource Semantic Parsing	Nathan Schucher, Siva Reddy, Harm de Vries
To Find Waldo You Need Contextual Cues: Debiasing Who's Waldo	Yiran Luo, Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, Chitta Baral
Towards Consistent Document-level Entity Linking: Joint Models for Entity Linking and Coreference Resolution	Klim Zaporozjets, Johannes Deleu, Yiwei Jiang, Thomas Demeester, Chris Develder
Towards Fair Evaluation of Dialogue State Tracking by Flexible Incorporation of Turn-level Performances	Suvodip Dey, Ramamohan Kummara, Maunendra Sankar Desarkar
Translate-Train Embracing Translationese Artifacts	Sicheng Yu, Qianru Sun, Hao Zhang, Jing Jiang
Triangular Transfer: Freezing the Pivot for Triangular Machine Translation	Meng Zhang, Liangyou Li, Qun Liu
(Un)solving Morphological Inflection: Lemma Overlap Artificially Inflates Models' Performance	Omer Goldman, David Guriel, Reut Tsarfaty

Understanding Game-Playing Agents with Natural Language Annotations	Nicholas Tomlin, Andre Wang He, Dan Klein
UniGDD: A Unified Generative Framework for Goal-Oriented Document-Grounded Dialogue	Chang Gao, Wenxuan Zhang, Wai Lam
Unsupervised multiple-choice question generation for out-of-domain Q&A fine-tuning	Guillaume Le Berre, Christophe Cerisara, Philippe Langlais, Guy Lapalme
Voxel-informed Language Grounding	Rodolfo Corona, Shizhan Zhu, Dan Klein, Trevor Darrell
When classifying grammatical role, BERT doesn't care about word order... except when it matters	Isabel Papadimitriou, Richard Futrell, Kyle Mahowald
When to Use Multi-Task Learning vs Intermediate Fine-Tuning for Pre-Trained Encoder Transfer Learning	Orion Weller, Kevin Seppi, Matt Gardner
WLASL-LEX: a Dataset for Recognising Phonological Properties in American Sign Language	Federico Tavella, Viktor Schlegel, Marta Romeo, Aphrodite Galata, Angelo Cangelosi
XDBERT: Distilling Visual Information to BERT from Cross-Modal Systems to Improve Language Understanding	Chan-Jan Hsu, Hung-yi Lee, Yu Tsao
Zero-Shot Dependency Parsing with Worst-Case Aware Automated Curriculum Learning	Miryam de Lhoneux, Sheng Zhang, Anders Søgaard

System Demonstration Papers

Title	Authors
DoTAT: A Domain-oriented Text Annotation Tool	Yupian Lin, Tong Ruan, Ming Liang, Tingting Cai, Wen Du, Yi Wang
UKP-SQUARE: An Online Platform for Question Answering Research	Tim Baumgärtner, Kexin Wang, Rachneet Sachdeva, Gregor Geigle, Max Eichler, Clifton Poth, Hannah Sterz, Haritz Puerto, Leonardo F. R. Ribeiro, Jonas Pfeiffer, Nils Reimers, Gözde Şahin, Iryna Gurevych
VILMedic: a framework for research at the intersection of vision and language in medical AI	Jean-benoît Delbrouck, Khaled Saab, Maya Varma, Sabri Eyuboglu, Pierre Chambon, Jared Dunnmon, Juan Zambrano, Akshay Chaudhari, Curtis Langlotz
TextPruner: A Model Pruning Toolkit for Pre-Trained Language Models	Ziqing Yang, Yiming Cui, Zhigang Chen
AnnIE: An Annotation Platform for Constructing Complete Open Information Extraction Benchmark	Niklas Friedrich, Kiril Gashteovski, Mingying Yu, Bhushan Kotnis, Carolin Lawrence, Mathias Niepert, Goran Glavaš
AdapterHub Playground: Simple and Flexible Few-Shot Learning with Adapters	Tilman Beck, Bela Bohlender, Christina Viehmann, Vincent Hane, Yanik Adamson, Jaber Khuri, Jonas Brossmann, Jonas Pfeiffer, Iryna Gurevych
QiuNiu: A Chinese Lyrics Generation System with Passage-Level Input	Le Zhang, Rongsheng Zhang, Xiaoxi Mao, Yongzhu Chang

Automatic Gloss Dictionary for Sign Language Learners	Chenchen Xu, Dongxu Li, Hongdong Li, Hanna Suominen, Ben Swift
PromptSource: An Integrated Development Environment and Repository for Natural Language Prompts	Stephen Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-david, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Fries, Maged Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Dragomir Radev, Mike Tian-jian Jiang, Alexander Rush
OpenPrompt: An Open-source Framework for Prompt-learning	Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, Maosong Sun
Guided K-best Selection for Semantic Parsing Annotation	Anton Belyy, Chieh-yang Huang, Jacob Andreas, Emmanouil Antonios Platanios, Sam Thomson, Richard Shin, Subhro Roy, Aleksandr Nisnevich, Charles Chen, Benjamin Van Durme
Hard and Soft Evaluation of NLP models with BOOtSTrap SAMpling - BooStSa	Tommaso Fornaciari, Alexandra Uma, Massimo Poesio, Dirk Hovy
COVID-19 Claim Radar: A Structured Claim Extraction and Tracking System	Manling Li, Revanth Gangi Reddy, Ziqi Wang, Yishyuan Chiang, Tuan Lai, Pengfei Yu, Zixuan Zhang, Heng Ji
TS-ANNO: An Annotation Tool to Build, Annotate and Evaluate Text Simplification Corpora	Regina Stodden, Laura Kallmeyer
Language Diversity: Visible to Humans, Exploitable by Machines	Gábor Bella, Erdenebileg Byambadorj, Yamini Chandrashekar, Khuyagbaatar Batsuren, Danish Cheema, Fausto Giunchiglia
CogKGE: A Knowledge Graph Embedding Toolkit and Benchmark for Representing Multi-source and Heterogeneous Knowledge	Zhuoran Jin, Tianyi Men, Hongbang Yuan, Zhitao He, Dianbo Sui, Chenhao Wang, Zhipeng Xue, Yubo Chen, Jun Zhao
Dynatask: A Framework for Creating Dynamic AI Benchmark Tasks	Tristan Thrush, Kushal Tirumala, Anmol Gupta, Max Bartolo, Pedro Rodriguez, Tariq Kane, William Gaviria Rojas, Peter Mattson, Adina Williams, Douwe Kiela
DataLab: A Platform for Data Analysis and Intervention	Yang Xiao, Jinlan Fu, Weizhe Yuan, Vijay Viswanathan, Zhoumianze Liu, Yixin Liu, Graham Neubig, Pengfei Liu
Cue-bot: A Conversational Agent for Assistive Technology	Shachi H Kumar, Hsuan Su, Ramesh Manuvinakurike, Maximilian C Pinaroc, Sai Prasad, Saurav Sahay, Lama Nachman
M-SENA: An Integrated Platform for Multimodal Sentiment Analysis	Huisheng Mao, Ziqi Yuan, Hua Xu, Wenmeng Yu, Yihe Liu, Kai Gao
HOSMEL: A Hot-Swappable Modularized Entity Linking Toolkit for Chinese	Daniel Zhang-li, Jing Zhang, Jifan Yu, Xiaokang Zhang, Peng Zhang, Jie Tang, Juanzi Li
BMInf: An Efficient Toolkit for Big Model Inference and Tuning	Xu Han, Guoyang Zeng, Weilin Zhao, Zhiyuan Liu,

	Zhengyan Zhang, Jie Zhou, Jun Zhang, Jia Chao, Maosong Sun
MMEKG: Multi-modal Event Knowledge Graph towards Universal Representation across Modalities	Yubo Ma, Zehao Wang, Mukai Li, Yixin Cao, Meiqi Chen, Xinze Li, Wenqi Sun, Kunquan Deng, Kun Wang, Aixin Sun, Jing Shao
SocioFillmore: A Tool for Discovering Perspectives	Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, Malvina Nissim
TimeLMs: Diachronic Language Models from Twitter	Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, Jose Camacho-collados
Adaptor: Objective-Centric Adaptation Framework for Language Models	Michal Štefánik, Vít Novotný, Nikola Groverová, Petr Sojka
QuickGraph: A Rapid Annotation Tool for Knowledge Graph Extraction from Technical Text	Tyler Bikaun, Michael Stewart, Wei Liu

FINDINGS OF ACL 2022

Long Papers

Title	Authors
<i>A Few-Shot Semantic Parser for Wizard-of-Oz Dialogues with the Precise ThingTalk Representation</i>	Giovanni Campagna, Sina Semnani, Ryan Kearns, Lucas Jun Koba Sato, Silei Xu, Monica Lam
<i>A Graph Enhanced BERT Model for Event Prediction</i>	Li Du, Xiao Ding, Yue Zhang, ting liu, Bing Qin
<i>A Natural Diet: Towards Improving Naturalness of Machine Translation Output</i>	Markus Freitag, David Vilar, David Grangier, Colin Cherry, George Foster
<i>A Neural Pairwise Ranking Model for Readability Assessment</i>	Justin Lee, Sowmya Vajjala
<i>A Novel Framework Based on Medical Concept Driven Attention for Explainable Medical Code Prediction via External Knowledge</i>	Tao Wang, Linhai Zhang, Chenchen Ye, Junxi Liu, Deyu Zhou
<i>A Novel Perspective to Look At Attention: Bi-level Attention-based Explainable Topic Modeling for News Classification</i>	Dairui Liu, Derek Greene, Ruihai Dong
<i>A Sentence is Worth 128 Pseudo Tokens: A Semantic-Aware Contrastive Learning Framework for Sentence Embeddings</i>	Haochen Tan, Wei Shao, Han Wu, Ke Yang, Linqi Song
<i>A Simple Hash-Based Early Exiting Approach For Language Understanding and Generation</i>	Tianxiang Sun, Xiangyang Liu, Wei Zhu, Zhichao Geng, Lingling Wu, Yilong He, Yuan Ni, Guotong Xie, Xuanjing Huang, Xipeng Qiu
<i>A Slot Is Not Built in One Utterance: Spoken Language Dialogs with Sub-Slots</i>	Sai Zhang, Yuwei Hu, Yuchuan Wu, Jiaman Wu, Yongbin Li, Jian SUN, Caixia Yuan, Xiaojie Wang
<i>A Transformational Biencoder with In-Domain Negative Sampling for Zero-Shot Entity Linking</i>	Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, Xudong Liu

<i>AbductionRules: Training Transformers to Explain Unexpected Inputs</i>	Nathan Young, Qiming Bao, Joshua Ljudo Bensemann, Michael J. Witbrock
<i>Aligned Weight Regularizers for Pruning Pretrained Neural Networks</i>	James O' Neill, Sourav Dutta, Haytham Assem
<i>AMR-DA: Data Augmentation by Abstract Meaning Representation</i>	Ziyi Shou, Yuxin Jiang, Fangzhen Lin
<i>An Accurate Unsupervised Method for Joint Entity Alignment and Dangling Entity Detection</i>	Shengxuan Luo, Sheng Yu
<i>Analyzing Dynamic Adversarial Training Data in the Limit</i>	Eric Wallace, Adina Williams, Robin Jia, Douwe Kiela
<i>Answer Uncertainty and Unanswerability in Multiple-Choice Machine Reading Comprehension</i>	Vatsal Raina, Mark Gales
<i>ASCM: An Answer Space Clustered Prompting Method without Answer Engineering</i>	Zhen Wang, Yating Yang, Zhou Xi, Bo Ma, Lei Wang, Rui Dong, Azmat Anwar
<i>Assessing Multilingual Fairness in Pre-trained Multimodal Representations</i>	Jialu Wang, Yang Liu, Xin Eric Wang
<i>ASSIST: Towards Label Noise-Robust Dialogue State Tracking</i>	Fanghua Ye, Yue Feng, Emine Yilmaz
<i>Attention as Grounding: Exploring Textual and Cross-Modal Attention on Entities and Relations in Language-and-Vision Transformer</i>	Nikolai Ilinykh, Simon Dobnik
<i>Automatic Song Translation for Tonal Languages</i>	Fenfei Guo, Chen Zhang, Zhirui Zhang, Qixin He, Kejun Zhang, Jun Xie, Jordan Lee Boyd-Graber
<i>BBQ: A hand-built bias benchmark for question answering</i>	Alicia Vail Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, Samuel R. Bowman
<i>Better Quality Estimation for Low Resource Corpus Mining</i>	Muhammed Yusuf Kocuyigit, Jiho Lee, Derry Wijaya
<i>BiSyn-GAT+: Bi-Syntax Aware Graph Attention Network for Aspect-based Sentiment Analysis</i>	Shuo Liang, Wei Wei, Xian-Ling Mao, Fei Wang, Zhiyong He
<i>Breaking Down Multilingual Machine Translation</i>	Ting-Rui Chiang, Yi-Pei Chen, Yi-Ting Yeh, Graham Neubig
<i>Bridging Pre-trained Language Models and Hand-crafted Features for Unsupervised POS Tagging</i>	Houquan Zhou, Yang Li, Zhenghua Li, Min Zhang
<i>C3KG: A Chinese Commonsense Conversation Knowledge Graph</i>	Dawei Li, Yanran Li, Jiayi Zhang, Ke Li, Chen Wei, Jianwei Cui, Bin Wang
<i>Calibration of Machine Reading Systems at Scale</i>	Shehzaad Zuzar Dhuliawala, Leonard Adolphs, Rajarshi Das, Mrinmaya Sachan
<i>CaM-Gen: Causally Aware Metric-Guided Text Generation</i>	Navita Goyal, Roodram Paneri, Ayush Agarwal, Udit Kalani, Abhilasha Sancheti, Niyati Chhaya
<i>Capture Human Disagreement Distributions by Calibrated Networks for Natural Language Inference</i>	Yuxia Wang, Minghan Wang, Yimeng Chen, Shimin Tao, Jiaxin Guo, Chang Su, Min Zhang, Hao Yang
<i>Challenges to Open-Domain Constituency Parsing</i>	Sen Yang, Leyang Cui, Ruoxi Ning, Di Wu, Yue

	Zhang
<i>ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning</i>	Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, Enamul Hoque
<i>Chinese Synesthesia Detection: New Dataset and Models</i>	Xiaotong Jiang, Qingqing Zhao, YUNFEI LONG, Zhongqing Wang
<i>Classification without (Proper) Representation: Political Heterogeneity in Social Media and Its Implications for Classification and Behavioral Analysis</i>	Kenan Alkiek, Bohan Zhang, David Jurgens
<i>Co-training an Unsupervised Constituency Parser with Weak Supervision</i>	Nickil Maveli, Shay B Cohen
<i>Co-VQA : Answering by Interactive Sub Question Sequence</i>	Ruonan Wang, Yuxi Qian, Fangxiang Feng, Xiaojie Wang, Huixing Jiang
<i>CoCoLM: Complex Commonsense Enhanced Language Model with Discourse Relations</i>	Changlong Yu, Hongming Zhang, Yangqiu Song, Wilfred Ng
<i>Coloring the Blank Slate: Pre-training Imparts a Hierarchical Inductive Bias to Sequence-to-sequence Models</i>	Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang, Sebastian Schuster
<i>Combining Feature and Instance Attribution to Detect Artifacts</i>	Pouya Pezeshkpour, Sarthak Jain, Sameer Singh, Byron C Wallace
<i>Comparative Opinion Summarization via Collaborative Decoding</i>	Hayate Iso, Xiaolan Wang, Stefanos Angelidis, Yoshihiko Suhara
<i>Compilable Neural Code Generation with Compiler Feedback</i>	Xin Wang, Yasheng Wang, Yao Wan, Fei Mi, Yitong Li, Pingyi Zhou, Jin Liu, Hao Wu, Xin Jiang, Qun Liu
<i>Composing Structure-Aware Batches for Pairwise Sentence Classification</i>	Andreas Waldis, Tilman Beck, Iryna Gurevych
<i>Comprehensive Multi-Modal Interactions for Referring Image Segmentation</i>	Kanishk Jain, Vineet Gandhi
<i>Consistent Representation Learning for Continual Relation Extraction</i>	Kang Zhao, Hua Xu, Jiangong Yang, Kai Gao
<i>Constructing Open Cloze Tests Using Generation and Discrimination Capabilities of Transformers</i>	Mariano Felice, Shiva Taslimipoor, Paula Buttery
<i>Controllable Natural Language Generation with Contrastive Prefixes</i>	Jing Qian, Li Dong, Yelong Shen, Furu Wei, Weizhu Chen
<i>Controlled Text Generation Using Dictionary Prior in Variational Autoencoders</i>	Xianghong Fang, Jian Li, Lifeng Shang, Xin Jiang, Qun Liu, Dit-Yan Yeung
<i>Controlling the Focus of Pretrained Language Generation Models</i>	Jiabao Ji, Yoon Kim, James R. Glass, Tianxing He
<i>CRAFT: A Benchmark for Causal Reasoning About Forces and Interactions</i>	Tayfun Ates, M. Şamil Ateşoğlu, Çağatay Yiğit, Ilker Kesen, Mert Kobas, Erkut Erdem, Aykut Erdem, Tilbe Goksun, Deniz Yuret
<i>CRASpell: A Contextual Typo Robust Approach to Improve Chinese Spelling Correction</i>	Shulin Liu, Shengkang Song, Tianchi Yue, TAO YANG, Huihui Cai, TingHao YU, Shengli Sun
<i>Cross-domain Named Entity Recognition via Graph Matching</i>	Junhao Zheng, Haibin Chen, Qianli Ma

<i>Cross-lingual Inference with A Chinese Entailment Graph</i>	Tianyi Li, Sabine Weber, Mohammad Javad Hosseini, Liane Guillou, Mark Steedman
<i>Cross-Lingual UMLS Named Entity Linking using UMLS Dictionary Fine-Tuning</i>	Rina Galperin, Shachar Schnapp, Michael Elhadad
<i>Cross-Modal Cloze Task: A New Task to Brain-to-Word Decoding</i>	Shuxian Zou, Shaonan Wang, Jiajun Zhang, Chengqing Zong
<i>CrossAligner & Co: Zero-Shot Transfer Methods for Task-Oriented Cross-lingual Natural Language Understanding</i>	Milan Gritta, Ruoyu Hu, Ignacio Iacobacci
<i>CUE Vectors: Modular Training of Language Models Conditioned on Diverse Contextual Signals</i>	Scott Novotney, Sreeparna Mukherjee, Zeeshan Ahmed, Andreas Stolcke
<i>Cutting Down on Prompts and Parameters: Simple Few-Shot Learning with Language Models</i>	Robert L. Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh, Sebastian Riedel
<i>DaLC: Domain Adaptation Learning Curve Prediction for Neural Machine Translation</i>	Cheonbok Park, Hantae Kim, Ioan Calapodescu, Hyun Chang Cho, Vassilina Nikoulina
<i>DARER: Dual-task Temporal Relational Recurrent Reasoning Network for Joint Dialog Sentiment Classification and Act Recognition</i>	Bowen Xing, Ivor Tsang
<i>Data Augmentation and Learned Layer Aggregation for Improved Multilingual Language Understanding in Dialogue</i>	Evgeniia Razumovskaia, Ivan Vulić, Anna Korhonen
<i>Decomposed Meta-Learning for Few-Shot Named Entity Recognition</i>	Tingting Ma, Huiqiang Jiang, Qianhui Wu, Tiejun Zhao, Chin-Yew Lin
<i>Deep Reinforcement Learning for Entity Alignment</i>	Lingbing Guo, Yuqiang Han, Qiang Zhang, Huajun Chen
<i>DeepStruct: Pre-Training of Language Models for Structure Prediction</i>	Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, Dawn Song
<i>Detecting Various Types of Noise for Neural Machine Translation</i>	Christian Herold, Jan Rosendahl, Joris Vanvinckenroye, Hermann Ney
<i>Detection of Adversarial Examples in Text Classification: Benchmark and Baseline via Robust Density Estimation</i>	KiYoon Yoo, Jangho Kim, Jiho Jang, Nojun Kwak
<i>Detection, Disambiguation, Re-ranking: Autoregressive Entity Linking as a Multi-Task Problem</i>	Khalil Mrini, Shaoliang Nie, Jiatao Gu, Sinong Wang, Maziar Sanjabi, Hamed Firooz
<i>Dialogue Summaries as Dialogue States (DS2), Template-Guided Summarization for Few-shot Dialogue State Tracking</i>	Jamin Shin, Hangyeol Yu, Hyeongdon Moon, Andrea Madotto, Juneyoung Park
<i>Dict-BERT: Enhancing Language Model Pre-training with Dictionary</i>	Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng, Meng Jiang
<i>Dim Wohl Gat Tun: The Case for Linguistic Expertise in NLP for Under-Documented Languages</i>	Clarissa Forbes, Farhan Samir, Bruce Harold Oliver, Changbing Yang, Edith Coates, Garrett Nicolai, Miikka Silfverberg
<i>Discontinuous Constituency and BERT: A Case Study of Dutch</i>	Konstantinos Kogkalidis, Gijs Wijnholds
<i>Distinguishing Non-natural from Natural Adversarial Samples for More</i>	Jiayi Wang, Rongzhou Bao, Zhuosheng Zhang,

<i>Robust Pre-trained Language Model</i>	hai zhao
<i>Distributed NLI: Learning to Predict Human Opinion Distributions for Language Reasoning</i>	Xiang Zhou, Yixin Nie, Mohit Bansal
<i>Diversifying Content Generation for Commonsense Reasoning with Mixture of Knowledge Graph Experts</i>	Wenhao Yu, Chenguang Zhu, Lianhui Qin, Zhihan Zhang, Tong Zhao, Meng Jiang
<i>Divide and Conquer: Text Semantic Matching with Disentangled Keywords and Intents</i>	Yicheng Zou, Hongwei Liu, Tao Gui, Junzhe Wang, Qi Zhang, Meng Tang, Haixiang Li, daniell wang
<i>Do Pre-trained Models Benefit Knowledge Graph Completion? A Reliable Evaluation and a Reasonable Approach</i>	Xin Lv, Yankai Lin, Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Peng Li, Jie Zhou
<i>Document-Level Event Argument Extraction via Optimal Transport</i>	Amir Pouran Ben Veyseh, Minh Van Nguyen, Franck Dernoncourt, Bonan Min, Thien Huu Nguyen
<i>Document-Level Relation Extraction with Adaptive Focal Loss and Knowledge Distillation</i>	Qingyu Tan, Ruidan He, Lidong Bing, Hwee Tou Ng
<i>Domain Representative Keywords Selection: A Probabilistic Approach</i>	Pritom Saha Akash, Jie Huang, Kevin Chang, Yunyao Li, Lucian Popa, ChengXiang Zhai
<i>DS-TOD: Efficient Domain Specialization for Task-Oriented Dialog</i>	Chia-Chien Hung, Anne Lauscher, Simone Paolo Ponzetto, Goran Glavaš
<i>DU-VLG: Unifying Vision-and-Language Generation via Dual Sequence-to-Sequence Pre-training</i>	Luyang Huang, Guocheng Niu, Jiachen Liu, Xinyan Xiao, Hua Wu
<i>DuReader_vis: A Chinese Dataset for Open-domain Document Visual Question Answering</i>	Le Qi, Shangwen Lv, Hongyu Li, Jing Liu, Yu Zhang, Qiaoqiao She, Hua Wu, Haifeng Wang, Ting Liu
<i>E-KAR: A Benchmark for Rationalizing Natural Language Analogical Reasoning</i>	Jiangjie Chen, Rui Xu, Ziquan Fu, Wei Shi, Zhongqiao Li, Xinbo Zhang, Changzhi Sun, Lei Li, Yanghua Xiao, Hao Zhou
<i>ECO v1: Towards Event-Centric Opinion Mining</i>	Ruoxi Xu, Hongyu Lin, Meng Liao, Xianpei Han, Jin Xu, Wei Tan, Yingfei Sun, Le Sun
<i>ED2LM: Encoder-Decoder to Language Model for Faster Document Re-ranking Inference</i>	Kai Hui, Honglei Zhuang, Tao Chen, Zhen Qin, Jing Lu, Dara Bahri, Ji Ma, Jai Gupta, Cicero Nogueira dos Santos, Yi Tay, Donald Metzler
<i>Effective Unsupervised Constrained Text Generation based on Perturbed Masking</i>	Yingwen Fu, Wenjie Ou, Zhou Yu, Yue Lin
<i>Efficient Argument Structure Extraction with Transfer Learning and Active Learning</i>	Xinyu Hua, Lu Wang
<i>Efficient, Uncertainty-based Moderation of Neural Networks Text Classifiers</i>	Jakob Smedegaard Andersen, Walid Maalej
<i>Eider: Empowering Document-level Relation Extraction with Efficient Evidence Extraction and Inference-stage Fusion</i>	Yiqing Xie, Jiaming Shen, Sha Li, Yuning Mao, Jiawei Han
<i>ELLE: Efficient Lifelong Pre-training for Emerging Data</i>	Yujia Qin, Jiajie Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou

<i>EmoCaps: Emotion Capsule based Model for Conversational Emotion Recognition</i>	Zaijing Li, Fengxiao Tang, Ming Zhao, Yusen Zhu
<i>Enabling Multimodal Generation on CLIP via Vision-Language Knowledge Distillation</i>	Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, Pascale Fung
<i>EnCBP: A New Benchmark Dataset for Finer-Grained Cultural Background Prediction in English</i>	Weicheng Ma, Samiha Datta, Lili Wang, Soroush Vosoughi
<i>Encoding and Fusing Semantic Connection and Linguistic Evidence for Implicit Discourse Relation Recognition</i>	Wei Xiang, Bang Wang, Lu Dai, Yijun Mo
<i>End-to-End Segmentation-based News Summarization</i>	Yang Liu, Chenguang Zhu, Michael Zeng
<i>End-to-End Speech Translation for Code Switched Speech</i>	Orion Weller, Matthias Sperber, Telmo Pires, Hendra Setiawan, Christian Gollan, Dominic C Telaar, Matthias Paulik
<i>Enhancing Natural Language Representation with Large-Scale Out-of-Domain Commonsense</i>	Wanyun Cui, Xingran Chen
<i>Entropy-based Attention Regularization Frees Unintended Bias Mitigation from Lists</i>	Giuseppe Attanasio, Debora Nozza, Dirk Hovy, Elena Baralis
<i>Event Transition Planning for Open-ended Text Generation</i>	Qintong Li, Piji Li, Wei Bi, Zhaochun Ren, Yuxuan Lai, Lingpeng Kong
<i>Explaining Classes through Stable Word Attributions</i>	Samuel Rönqvist, AKI-Juhani Kyröläinen, Amanda Myntti, Filip Ginter, Veronika Laippala
<i>Exploring the Capacity of a Large-scale Masked Language Model to Recognize Grammatical Errors</i>	Ryo Nagata, Manabu Kimura, Kazuaki Hanawa
<i>Exploring the Impact of Negative Samples of Contrastive Learning: A Case Study of Sentence Embedding</i>	Rui Cao, Yihao Wang, Yuxin Liang, Ling Gao, Jie Zheng, Jie Ren, Zheng Wang
<i>Extract-Select: A Span Selection Framework for Nested Named Entity Recognition with Generative Adversarial Training</i>	Peixin Huang, Xiang Zhao, Minghao Hu, Yang Fang, Xinyi Li, Weidong Xiao
<i>Extracting Latent Steering Vectors from Pretrained Language Models</i>	Nishant Subramani, Nivedita Suresh, Matthew E Peters
<i>Extracting Person Names from User Generated Text: Named-Entity Recognition for Combating Human Trafficking</i>	Yifei Li, Pratheeksha Nair, Kellin Pelrine, Reihaneh Rabbany
<i>Fact-Tree Reasoning for N-ary Question Answering over Knowledge Graphs</i>	YAO ZHANG, Peiyao Li, Hongru Liang, Adam Jatowt, Zhenglu Yang
<i>Fast Nearest Neighbor Machine Translation</i>	Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang, Jiwei Li
<i>Finding the Dominant Winning Ticket in Pre-Trained Language Models</i>	Zhuocheng Gong, Di He, Yelong Shen, Tie-Yan Liu, Weizhu Chen, Dongyan Zhao, Ji-Rong Wen, Rui Yan
<i>Focus on the Action: Learning to Highlight and Summarize Jointly for Email To-Do Items Summarization</i>	Kexun Zhang, Jiaao Chen, Diyi Yang
<i>From BERT's Point of View: Revealing the Prevailing Contextual</i>	Carolyn M. Schuster, Simon Hegelich

<i>Differences</i>	
<i>From Stance to Concern: Adaptation of Propositional Analysis to New Tasks and Domains</i>	Brodie Mather, Bonnie J Dorr, Adam Dalton, William de Beaumont, Owen Rambow, Sonja M. Schmer-Galunder
<i>Fusing Heterogeneous Factors with Triaffine Mechanism for Nested Named Entity Recognition</i>	Zheng Yuan, Chuanqi Tan, Songfang Huang, Fei Huang
<i>Gaussian Multi-head Attention for Simultaneous Machine Translation</i>	Shaolei Zhang, Yang Feng
<i>GCPG: A General Framework for Controllable Paraphrase Generation</i>	Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Haibo Zhang, Xue Zhao, Wenqing Yao, Boxing Chen
<i>Generalized but not Robust? Comparing the Effects of Data Modification Methods on Out-of-Domain Generalization and Adversarial Robustness</i>	Tejas Gokhale, Swaroop Mishra, Man Luo, Bhavdeep Singh Sachdeva, Chitta Baral
<i>Going "Deeper": Structured Sememe Prediction via Transformer with Tree Attention</i>	Yining Ye, Fanchao Qi, Zhiyuan Liu, Maosong Sun
<i>Good Night at 4 pm?! Time Expressions in Different Cultures</i>	Vered Shwartz
<i>Graph Neural Networks for Multiparallel Word Alignment</i>	Ayyoob Imani, Lütfi Kerem Senel, Masoud Jalili Sabet, François Yvon, Hinrich Schuetze
<i>Graph Refinement for Coreference Resolution</i>	Lesly Miculicich, James Henderson
<i>GRS: Combining Generation and Revision in Unsupervised Sentence Simplification</i>	Mohammad Dehghan, Dhruv Kumar, Lukasz Golab
<i>Hey AI, Can You Solve Complex Tasks by Talking to Agents?</i>	Tushar Khot, Kyle Richardson, Daniel Khoshabi, Ashish Sabharwal
<i>HiCLRE: A Hierarchical Contrastive Learning Framework for Distantly Supervised Relation Extraction</i>	Dongyang Li, Taolin Zhang, Nan Hu, Chengyu Wang, Xiaofeng He
<i>HIE-SQL: History Information Enhanced Network for Context-Dependent Text-to-SQL Semantic Parsing</i>	Yanzhao Zheng, Haibin Wang, Dong Baohua, Xingjun Wang, Changshan Li
<i>Hierarchical Recurrent Aggregative Generation for Few-Shot NLG</i>	Giulio Zhou, Gerasimos Lampouras, Ignacio Iacobacci
<i>HiStruct+: Improving Extractive Text Summarization with Hierarchical Structure Information</i>	Qian Ruan, Malte Ostendorff, Georg Rehm
<i>HLDC: Hindi Legal Documents Corpus</i>	Arnav Kapoor, Mudit Dhawan, Anmol Goel, Arjun T H, Akshala Bhatnagar, Vibhu Agrawal, Amul Agrawal, Arnab Bhattacharya, Ponnurangam Kumaraguru, Ashutosh Modi
<i>How Can Cross-lingual Knowledge Contribute Better to Fine-Grained Entity Typing?</i>	Hailong Jin, Tiansi Dong, Lei Hou, Juanzi Li, Hui Chen, Zelin Dai, Qu Yincen
<i>How Pre-trained Language Models Capture Factual Knowledge? A Causal-Inspired Analysis</i>	Shaobo Li, Xiaoguang Li, Lifeng Shang, Zhenhua Dong, Chengjie Sun, Bingquan Liu, Zhenzhou Ji, Xin Jiang, Qun Liu
<i>Human Language Modeling</i>	Nikita Soni, Matthew Matero, Niranjana

	Balasubramanian, H. Schwartz
<i>HybridDialogue: An Information-Seeking Dialogue Dataset Grounded on Tabular and Textual Data</i>	Kai Nakamura, Sharon Levy, Yi-Lin Tuan, Wenhui Chen, William Yang Wang
<i>Implicit Relation Linking for Question Answering over Knowledge Graph</i>	Yao Zhao, Jiacheng Huang, Wei Hu, Qijin Chen, XiaoXia Qiu, Chengfu Huo, Weijun Ren
<i>Improved Multi-label Classification under Temporal Concept Drift: Rethinking Group-Robust Algorithms in a Label-Wise Setting</i>	Ilias Chalkidis, Anders Søgaard
<i>Improving Candidate Retrieval with Entity Profile Generation for Wikidata Entity Linking</i>	Tuan Lai, Heng Ji, ChengXiang Zhai
<i>Improving Chinese Grammatical Error Detection via Data augmentation by Conditional Error Generation</i>	Tianchi Yue, Shulin Liu, Huihui Cai, TAO YANG, Shengkang Song, TingHao YU
<i>Improving Controllable Text Generation with Position-Aware Weighted Decoding</i>	Yuxuan Gu, Xiaocheng Feng, Sicheng Ma, Jiaming Wu, Heng Gong, Bing Qin
<i>Improving Neural Political Statement Classification with Class Hierarchical Information</i>	Erenay Dayanik, Andre Blessing, Nico Blokker, Sebastian Haunss, Jonas Kuhn, Gabriella Lapesa, Sebastian Pado
<i>Improving Relation Extraction through Syntax-induced Pre-training with Dependency Masking</i>	Yuanhe Tian, Yan Song, Fei Xia
<i>Improving Robustness of Language Models from a Geometry-aware Perspective</i>	Bin Zhu, Zhaoquan Gu, Le Wang, Jinyin Chen, Qi Xuan
<i>Improving the Adversarial Robustness of NLP Models by Information Bottleneck</i>	Cenyuan Zhang, Xiang Zhou, Yixin Wan, Xiaoqing Zheng, Kai-Wei Chang, Cho-Jui Hsieh
<i>Incorporating Dynamic Semantics into Pre-Trained Language Model for Aspect-based Sentiment Analysis</i>	Kai Zhang, Kun Zhang, Mengdi Zhang, Hongke Zhao, Qi Liu, Wei Wu, Enhong Chen
<i>IndicBART: A Pre-trained Model for Indic Natural Language Generation</i>	Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh M Khapra, Pratyush Kumar
<i>Input-specific Attention Subnetworks for Adversarial Detection</i>	Emil Biju, Anirudh Sriram, Pratyush Kumar, Mitesh M Khapra
<i>Interpretable Research Replication Prediction via Variational Contextual Consistency Sentence Masking</i>	Tianyi Luo, Rui Meng, Xin Eric Wang, Yang Liu
<i>Interpreting the Robustness of Neural NLP Models to Textual Perturbations</i>	Yunxiang Zhang, Liangming Pan, Samson Tan, Min-Yen Kan
<i>Inverse is Better! Fast and Accurate Prompt for Few-shot Slot Tagging</i>	Yutai Hou, Cheng Chen, Xianzhen Luo, Bohan Li, Wanxiang Che
<i>IsoScore: Measuring the Uniformity of Embedding Space Utilization</i>	William Rudman, Nate Gillman, Taylor Rayne, Carsten Eickhoff
<i>Knowledge Graph Embedding by Adaptive Limit Scoring Loss Using Dynamic Weighting Strategy</i>	Jinfa Yang, Xianghua Ying, Yongjie Shi, Xin Tong, Ruibin Wang, Taiyan Chen, Bowei Xing
<i>KSAM: Infusing Multi-Source Knowledge into Dialogue Generation via</i>	Sixing Wu, Ying Li, Dawei Zhang, Zhonghai Wu

<i>Knowledge Source Aware Multi-Head Decoding</i>	
<i>Label Semantics for Few Shot Named Entity Recognition</i>	Jie Ma, Miguel Ballesteros, Srikanth Doss, RISHITA ANUBHAI, Sunil Mallya, Yaser Al-Onaizan, Dan Roth
<i>Lacking the Embedding of a Word? Look it up into a Traditional Dictionary</i>	Elena Sofia Ruzzetti, Leonardo Ranaldi, Michele Mastromattei, Francesca Fallucchi, Noemi Scarpato, Fabio Massimo Zanzotto
<i>LaPraDoR: Unsupervised Pretrained Dense Retriever for Zero-Shot Text Retrieval</i>	Canwen Xu, Daya Guo, Nan Duan, Julian McAuley
<i>Learn and Review: Enhancing Continual Named Entity Recognition via Reviewing Synthetic Samples</i>	Yu Xia, Quan Wang, Yajuan Lyu, Yong Zhu, Wenhao Wu, Sujian Li, Dai Dai
<i>Learning Adaptive Axis Attentions in Fine-tuning: Beyond Fixed Sparse Attention Patterns</i>	Zihan Wang, Jiuxiang Gu, Jason Kuen, Handong Zhao, Vlad I Morariu, Ruiyi Zhang, Ani Nenkova, Tong Sun, Jingbo Shang
<i>Learning and Evaluating Character Representations in Novels</i>	Naoya Inoue, Charuta Pethe, Allen Kim, Steven Skiena
<i>Learning Bias-reduced Word Embeddings Using Dictionary Definitions</i>	Haozhe An, Xiaojiang Liu, Donald Zhang
<i>Learning from Missing Relations: Contrastive Learning with Commonsense Knowledge Graphs for Commonsense Inference</i>	Yong-Ho Jung, Jun-Hyung Park, Joon-Young Choi, Mingyu Lee, Junho Kim, Kang-Min Kim, SangKeun Lee
<i>Learning Reasoning Patterns for Relational Triple Extraction with Mutual Generation of Text and Graph</i>	Yubo Chen, Yunqi Zhang, Yongfeng Huang
<i>Learning to Describe Solutions for Bug Reports Based on Developer Discussions</i>	Sheena Panthaplackel, Junyi Jessy Li, Milos Gligoric, Ray Mooney
<i>Learning to Robustly Aggregate Labeling Functions for Semi-supervised Data Programming</i>	Ayush Maheshwari, Krishnateja Killamsetty, Ganesh Ramakrishnan, Rishabh K Iyer, Marina Danilevsky, Lucian Popa
<i>LEVEN: A Large-Scale Chinese Legal Event Detection Dataset</i>	Feng Yao, Chaojun Xiao, Xiaozhi Wang, Zhiyuan Liu, Lei Hou, Cunchao Tu, Juanzi Li, Yun Liu, Weixing Shen, Maosong Sun
<i>Leveraging Knowledge in Multilingual Commonsense Reasoning</i>	Yuwei Fang, Shuohang Wang, Yichong Xu, Ruochen Xu, Siqi Sun, Chenguang Zhu, Michael Zeng
<i>Listening to Affected Communities to Define Extreme Speech: Dataset and Experiments</i>	Antonis Maronikolakis, Axel Wisioerek, Leah Nann, Haris Jabbar, Sahana Udupa, Hinrich Schuetze
<i>Local Structure Matters Most: Perturbation Study in NLU</i>	Louis Clouatre, Prasanna Parthasarathi, Amal Zouaq, Sarath Chandar
<i>Logic-Driven Context Extension and Data Augmentation for Logical Reasoning of Text</i>	Siyuan Wang, Wanjun Zhong, Duyu Tang, zhongyu wei, Zhihao Fan, Daxin Jiang, Ming Zhou, Nan Duan
<i>Long Time No See! Open-Domain Conversation with Long-Term Persona Memory</i>	Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang, Shihang Wang

<i>MDCSpell: A Multi-task Detector-Corrector Framework for Chinese Spelling Correction</i>	Chenxi Zhu, Ziqiang Ying, Boyu Zhang, Feng Mao
<i>MDERank: A Masked Document Embedding Rank Approach for Unsupervised Keyphrase Extraction</i>	Linhan Zhang, Qian Chen, Wen Wang, Chong Deng, ShiLiang Zhang, Bing Li, Wei Wang, Xin Cao
<i>Measuring the Language of Self-Disclosure across Corpora</i>	Ann-Katrin Reuel, Sebastian Peralta, João Sedoc, Garrick Sherman, Lyle Ungar
<i>MERIt: Meta-Path Guided Contrastive Learning for Logical Reasoning</i>	Fangkai Jiao, Yangyang Guo, Xuemeng Song, Liqiang Nie
<i>Meta-X_NLG: A Meta-Learning Approach Based on Language Clustering for Zero-Shot Cross-Lingual Transfer and Generation</i>	Kaushal Kumar Maurya, Maunendra Sankar Desarkar
<i>Metadata Shaping: A Simple Approach for Knowledge-Enhanced Language Models</i>	Simran Arora, Sen Wu, Enci Liu, Christopher Re
<i>MetaWeighting: Learning to Weight Tasks in Multi-Task Learning</i>	Yuren Mao, Zekai Wang, Weiwei Liu, Xuemin Lin, Pengtao Xie
<i>MIMICause: Representation and automatic extraction of causal relation types from clinical notes</i>	Vivek Khetan, Md Imbesat Hassan Rizvi, Jessica Huber, Paige Bartusiak, Bogdan Eugen Sacaleanu, Andrew Fano
<i>MINER: Multi-Interest Matching Network for News Recommendation</i>	Jian Li, Jieming Zhu, Qiwei Bi, Guohao Cai, Lifeng Shang, Zhenhua Dong, Xin Jiang, Qun Liu
<i>Mitigating Gender Bias in Distilled Language Models via Counterfactual Role Reversal</i>	Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv Verma, Yada Pruksachatkun, Satyapriya Krishna, Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg, Aram Galstyan
<i>Mitigating the Inconsistency Between Word Saliency and Model Confidence with Pathological Contrastive Training</i>	Pengwei Zhan, Yang Wu, Shaolei Zhou, Yunjian Zhang, Liming Wang
<i>Modality-specific Learning Rates for Effective Multimodal Additive Late-fusion</i>	Yiqun Yao, Rada Mihalcea
<i>Modeling Intensification for Sign Language Generation: A Computational Approach</i>	Mert Inan, Yang Zhong, Sabit Hassan, Lorna Quandt, Malihe Alikhani
<i>Modular Domain Adaptation</i>	Junshen K Chen, Dallas Card, Dan Jurafsky
<i>MoEfication: Transformer Feed-forward Layers are Mixtures of Experts</i>	Zhengyan Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou
<i>More Than Words: Collocation Retokenization for Latent Dirichlet Allocation Models</i>	Jin Cheevaprawatdomrong, Alexandra Schofield, Attapol Rutherford
<i>Morphological Processing of Low-Resource Languages: Where We Are and What's Next</i>	Adam Wiemerslage, Miikka Silfverberg, Changbing Yang, Arya D. McCarthy, Garrett Nicolai, Eliana Colunga, Katharina Kann
<i>Morphosyntactic Tagging with Pre-trained Language Models for Arabic and its Dialects</i>	Go Inoue, Salam Khalifa, Nizar Habash
<i>MReD: A Meta-Review Dataset for Structure-Controllable Text Generation</i>	Chenhui Shen, Liying Cheng, Ran Zhou, Lidong

	Bing, Yang You, Luo Si
<i>Mukayese: Turkish NLP Strikes Back</i>	Ali Safaya, Emirhan Kurtuluş, Arda Goktogan, Deniz Yuret
<i>Multi-Granularity Semantic Aware Graph Model for Reducing Position Bias in Emotion Cause Pair Extraction</i>	Yinan Bao, Qianwen Ma, Lingwei Wei, Wei Zhou, Songlin Hu
<i>Multi-Scale Distribution Deep Variational Autoencoder for Explanation Generation</i>	ZeFeng Cai, Linlin Wang, Gerard de Melo, Fei Sun, Liang He
<i>Multi-Stage Prompting for Knowledgeable Dialogue Generation</i>	Zihan Liu, Mostofa Patwary, Ryan Prenger, Shrimai Prabhumoye, Wei Ping, Mohammad Shoeybi, Bryan Catanzaro
<i>Multi-task Learning for Paraphrase Generation With Keyword and Part-of-Speech Reconstruction</i>	Xuhang Xie, Xuesong Lu, Bei Chen
<i>N-Shot Learning for Augmenting Task-Oriented Dialogue State Tracking</i>	Ibrahim Taha Aksu, Zhengyuan Liu, Min-Yen Kan, Nancy F. Chen
<i>NEWS: A Corpus for News Topic-Focused Summarization</i>	Seyed Ali Bahrainian, Sheridan Feucht, Carsten Eickhoff
<i>OCR Improves Machine Translation for Low-Resource Languages</i>	Oana Ignat, Jean Maillard, Vishrav Chaudhary, Francisco Guzmán
<i>On Controlling Fallback Responses for Grounded Dialogue Generation</i>	Hongyuan Lu, Wai Lam, Hong Cheng, Helen M. Meng
<i>On the data requirements of probing</i>	Zining Zhu, Jixuan Wang, Bai Li, Frank Rudzicz
<i>On the Importance of Data Size in Probing Fine-tuned Models</i>	Houman Mehrfarin, Sara Rajaei, Mohammad Taher Pilehvar
<i>On the Safety of Conversational Models: Taxonomy, Dataset, and Benchmark</i>	Hao Sun, Guangxuan Xu, Jiawen Deng, Jiale Cheng, Chujie Zheng, Hao Zhou, Nanyun Peng, Xiaoyan Zhu, Minlie Huang
<i>One Agent To Rule Them All: Towards Multi-agent Conversational AI</i>	Christopher Clarke, Joseph Joshua Peper, Karthik Krishnamurthy, Walter Talamonti, Kevin Leach, Walter Lasecki, Yiping Kang, Lingjia Tang, Jason Mars
<i>OneAligner: Zero-shot Cross-lingual Transfer with One Rich-Resource Language Pair for Low-Resource Sentence Retrieval</i>	Tong Niu, Kazuma Hashimoto, Yingbo Zhou, Caiming Xiong
<i>Open Relation Modeling: Learning to Define Relations between Entities</i>	Jie Huang, Kevin Chang, Jinjun Xiong, Wen-mei Hwu
<i>Open Vocabulary Extreme Classification Using Generative Models</i>	Daniel Simig, Fabio Petroni, Pouya Yanki, Kashyap Popat, Christina Du, Sebastian Riedel, Majid Yazdani
<i>Perturbations in the Wild: Leveraging Human-Written Text Perturbations for Realistic Adversarial Attack and Defense</i>	Thai Le, Jooyoung Lee, Kevin Yen, Yifan Hu, Dongwon Lee
<i>Phoneme transcription of endangered languages: an evaluation of recent ASR architectures in the single speaker scenario</i>	Gilles Boulianne

<i>Platt-Bin: Efficient Posterior Calibrated Training for NLP Classifiers</i>	Rishabh Singh, Shirin Goshtasbpour
<i>Plug-and-Play Adaptation for Continuously-updated QA</i>	Kyungjae Lee, Wookje Han, seung-won hwang, Hwaran Lee, Joonsuk Park, Sang-Woo Lee
<i>Prior Knowledge and Memory Enriched Transformer for Sign Language Translation</i>	Tao Jin, Zhou Zhao, Meng Zhang, Xingshan Zeng
<i>Probing BERT's priors with serial reproduction chains</i>	Takateru Yamakoshi, Thomas L. Griffiths, Robert D. Hawkins
<i>Probing Factually Grounded Content Transfer with Factual Ablation</i>	Peter West, Chris Quirk, Michel Galley, Yejin Choi
<i>Probing Multilingual Cognate Prediction Models</i>	Clémentine Fourier, Benoît Sagot
<i>Prompt-Driven Neural Machine Translation</i>	Yafu Li, Yongjing Yin, Jing Li, Yue Zhang
<i>Query and Extract: Refining Event Extraction as Type-oriented Binary Decoding</i>	Sijia Wang, Mo Yu, Shiyu Chang, Lichao Sun, Lifu Huang
<i>Question Answering Infused Pre-training of General-Purpose Contextualized Representations</i>	Robin Jia, Mike Lewis, Luke Zettlemoyer
<i>Question Generation for Reading Comprehension Assessment by Modeling How and What to Ask</i>	Bilal Ghanem, Lauren Lutz Coleman, Julia Rivard Dexter, Spencer McIntosh von der Ohe, Alona Fyshe
<i>Ranking-Constrained Learning with Rationales for Text Classification</i>	Juanyan Wang, Manali Sharma, Mustafa Bilgic
<i>Read before Generate! Faithful Long Form Question Answering with Machine Reading</i>	Dan SU, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu, Pascale Fung
<i>Reinforced Cross-modal Alignment for Radiology Report Generation</i>	Han Qin, Yan Song
<i>RelationPrompt: Leveraging Prompts to Generate Synthetic Data for Zero-Shot Relation Triplet Extraction</i>	Yew Ken Chia, Lidong Bing, Soujanya Poria, Luo Si
<i>Rethinking Document-level Neural Machine Translation</i>	Zewei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, Lei Li
<i>Rethinking Offensive Text Detection as a Multi-Hop Reasoning Problem</i>	Qiang Zhang, Jason Naradowsky, Yusuke Miyao
<i>Revisiting Automatic Evaluation of Extractive Summarization Task: Can We Do Better than ROUGE?</i>	Mousumi Akter, Naman Bansal, Shubhra Kanti Karmaker
<i>Seeking Patterns, Not just Memorizing Procedures: Contrastive Learning for Solving Math Word Problems</i>	Zhongli Li, Wenxuan Zhang, Chao Yan, Qingyu Zhou, Chao Li, Hongzhi Liu, Yunbo Cao
<i>Semantically Distributed Robust Optimization for Vision-and-Language Inference</i>	Tejas Gokhale, Abhishek Chaudhary, Pratyay Banerjee, Chitta Baral, Yezhou Yang
<i>Sememe Prediction for BabelNet Synsets using Multilingual and Multimodal Information</i>	Fanchao Qi, Chuancheng Lv, Zhiyuan Liu, Xiaojun Meng, Maosong Sun, Hai-Tao Zheng
<i>Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models</i>	Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, Yinfei Yang

<i>Sentiment Word Aware Multimodal Refinement for Multimodal Sentiment Analysis with ASR Errors</i>	Yang Wu, Yanyan Zhao, Hao Yang, Song Chen, Bing Qin, Xiaohuan Cao, Wenting Zhao
<i>Seq2Path: Generating Sentiment Tuples as Paths of a Tree</i>	Yue Mao, Yi Shen, Jingchao Yang, Xiaoying Zhu, Longjun Cai
<i>Should We Trust This Summary? Bayesian Abstractive Summarization to The Rescue</i>	Alexios Gidiotis, Grigorios Tsoumakas
<i>Sibylvariant Transformations for Robust Text Classification</i>	Fabrice Y Harel-Canada, Muhammad Ali Gulzar, Nanyun Peng, Miryung Kim
<i>Square One Bias in NLP: Towards a Multi-Dimensional Exploration of the Research Manifold</i>	Sebastian Ruder, Ivan Vulić, Anders Søgaard
<i>Structural Supervision for Word Alignment and Machine Translation</i>	Lei Li, Kai Fan, Hongjia Li, Chun Yuan
<i>Suum Cuique: Studying Bias in Taboo Detection with a Community Perspective</i>	Osama Khalid, Jonathan Ruser, Padmini Srinivasan
<i>Synchronous Refinement for Neural Machine Translation</i>	Kehai Chen, Masao Utiyama, Eiichiro Sumita, Rui Wang, Min Zhang
<i>Syntax-guided Contrastive Learning for Pre-trained Language Model</i>	Shuai Zhang, Wang Lijie, Xinyan Xiao, Hua Wu
<i>Systematicity, Compositionality and Transitivity of Deep NLP Models: a Metamorphic Testing Perspective</i>	Edoardo Manino, Julia Rozanova, Danilo Carvalho, Andre Freitas, Lucas Carvalho Cordeiro
<i>TABi: Type-Aware Bi-Encoders for Open-Domain Entity Retrieval</i>	Megan Leszczynski, Daniel Y Fu, Mayee F Chen, Christopher Re
<i>Table-based Fact Verification with Self-adaptive Mixture of Experts</i>	Yuxuan Zhou, Xien Liu, Kaiyin Zhou, Ji Wu
<i>Task-guided Disentangled Tuning for Pretrained Language Models</i>	Jiali Zeng, Yufan Jiang, Shuangzhi Wu, Yongjing Yin, Mu Li
<i>TegTok: Augmenting Text Generation via Task-specific and Open-world Knowledge</i>	Chao-Hong Tan, Jia-Chen Gu, Chongyang Tao, Zhen-Hua Ling, Can Xu, Huang Hu, Xiubo Geng, Daxin Jiang
<i>Thai Nested Named Entity Recognition Corpus</i>	Weerayut Buaphet, Can Udomcharoenchaikit, Peerat Limkonchotiwat, Attapol Rutherford, Sarana Nutanong
<i>The Change that Matters in Discourse Parsing: Estimating the Impact of Domain Shift on Parser Error</i>	Katherine Atwell, Anthony Sicilia, Seong Jae Hwang, Malihe Alikhani
<i>The impact of lexical and grammatical processing on generating code from natural language</i>	Nathanaël Beau, Benoit Crabbé
<i>The Inefficiency of Language Models in Scholarly Retrieval: An Experimental Walk-through</i>	Shruti Singh, Mayank Singh
<i>The Past Mistake is the Future Wisdom: Error-driven Contrastive Probability Optimization for Chinese Spell Checking</i>	Yinghui Li, Qingyu Zhou, Yangning Li, Zhongli Li, Ruiyang Liu, Rongyi Sun, Zizhen Wang, Chao Li, Yunbo Cao, Hai-Tao Zheng
<i>THE-X: Privacy-Preserving Transformer Inference with Homomorphic Encryption</i>	Tianyu Chen, Hangbo Bao, Shaohan Huang, Li Dong, Binxing Jiao, Daxin Jiang, Haoyi Zhou,

	Jianxin Li, Furu Wei
<i>To be or not to be an Integer? Encoding Variables for Mathematical Text</i>	Deborah Ferreira, Mokanarangan Thayaparan, Marco Valentino, Julia Rozanova, Andre Freitas
<i>Toward More Meaningful Resources for Lower-resourced Languages</i>	Constantine Lignos, Nolan Holley, Chester Palen-Michel, Jonne Sälevä
<i>Towards Adversarially Robust Text Classifiers by Learning to Reweight Clean Examples</i>	Jianhan Xu, Cenyuan Zhang, Xiaoqing Zheng, Linyang Li, Cho-Jui Hsieh, Kai-Wei Chang, Xuanjing Huang
<i>Towards Collaborative Neural-Symbolic Graph Semantic Parsing via Uncertainty</i>	Zi Lin, Jeremiah Zhe Liu, Jingbo Shang
<i>Towards Few-shot Entity Recognition in Document Images: A Label-aware Sequence-to-Sequence Framework</i>	Zilong Wang, Jingbo Shang
<i>Towards Large-Scale Interpretable Knowledge Graph Reasoning for Dialogue Systems</i>	Yi-Lin Tuan, Sajjad Beygi, Maryam Fazel-Zarandi, Qiaozhi Gao, Alessandra Cervone, William Yang Wang
<i>Towards Transparent Interactive Semantic Parsing via Step-by-Step Correction</i>	Lingbo Mo, Ashley Lewis, Huan Sun, Michael White
<i>Towards Unifying the Label Space for Aspect- and Sentence-based Sentiment Analysis</i>	Yiming Zhang, Min Zhang, Sai Wu, Junbo Zhao
<i>Training Dynamics for Text Summarization Models</i>	Tanya Goyal, Jiacheng Xu, Junyi Jessy Li, Greg Durrett
<i>Training Text-to-Text Transformers with Privacy Guarantees</i>	Natalia Ponomareva, Jasmijn Bastings, Sergei Vassilvitskii
<i>Translation Error Detection as Rationale Extraction</i>	Marina Fomicheva, Lucia Specia, Nikolaos Aletras
<i>Type-Driven Multi-Turn Corrections for Grammatical Error Correction</i>	Shaopeng Lai, Qingyu Zhou, Jiali Zeng, Zhongli Li, Chao Li, Yunbo Cao, Jinsong Su
<i>UNIMO-2: End-to-End Unified Vision-Language Grounded Learning</i>	Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, Haifeng Wang
<i>Unsupervised Natural Language Inference Using PHL Triplet Generation</i>	Neeraj Varshney, Pratyay Banerjee, Tejas Gokhale, Chitta Baral
<i>Using Interactive Feedback to Improve the Accuracy and Explainability of Question Answering Systems Post-Deployment</i>	Zichao Li, Prakhar Sharma, Xing Han Lu, Jackie CK Cheung, Siva Reddy
<i>Using NLP to quantify the environmental cost and diversity benefits of in-person NLP conferences</i>	Piotr Przybyła, Matthew Shardlow
<i>Using Pre-Trained Language Models for Producing Counter Narratives Against Hate Speech: a Comparative Study</i>	Serra Sinem Tekiroglu, Helena Bonaldi, Margherita Fanton, Marco Guerini
<i>Virtual Augmentation Supported Contrastive Learning of Sentence Representations</i>	Dejiao Zhang, Wei Xiao, Henghui Zhu, Xiaofei Ma, Andrew Arnold
<i>VISITRON: Visual Semantics-Aligned Interactively Trained Object-Navigator</i>	Ayush Shrivastava, Karthik Gopalakrishnan, Yang Liu, Robinson Piramuthu, gokhan tur, Devi Parikh,

	Dilek Hakkani-Tur
<i>Visualizing the Relationship Between Encoded Linguistic Information and Task Performance</i>	Jiannan Xiang, Huayang Li, Defu Lian, Guoping Huang, Taro Watanabe, Lemao Liu
<i>Weighted self Distillation for Chinese word segmentation</i>	Rian He, Shubin Cai, Zhong Ming, Jiale Zhang
<i>What does it take to bake a cake? The RecipeRef corpus and anaphora resolution in procedural text</i>	Biaoyan Fang, Timothy Baldwin, Karin Verspoor
<i>What to Learn, and How: Toward Effective Learning from Rationales</i>	Samuel Carton, Surya Kanoria, Chenhao Tan
<i>What Works and Doesn't Work, A Deep Decoder for Neural Machine Translation</i>	Zuchao Li, Yiran Wang, Masao Utiyama, Eiichiro Sumita, hai zhao, Taro Watanabe
<i>When Chosen Wisely, More Data Is What You Need: A Universal Sample-Efficient Strategy For Data Augmentation</i>	Ehsan Kamaloo, Mehdi Rezagholizadeh, Ali Ghodsi
<i>Why don't people use character-level machine translation?</i>	Jindřich Libovický, Helmut Schmid, Alexander Fraser
<i>Why Exposure Bias Matters: An Imitation Learning Perspective of Error Accumulation in Language Generation</i>	Kushal Arora, Layla El Asri, Hareesh Bahuleyan, Jackie CK Cheung
<i>Word Segmentation by Separation Inference for East Asian Languages</i>	Yu Tong, Jingzhi Guo, Jizhe Zhou, Ge Chen, Guokai Zheng
<i>XFUND: A Benchmark Dataset for Multilingual Visually Rich Form Understanding</i>	Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Furu Wei
<i>xGQA: Cross-Lingual Visual Question Answering</i>	Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin O. Steitz, Stefan Roth, Ivan Vulić, Iryna Gurevych
<i>Your fairness may vary: Pretrained language model fairness in toxic text classification</i>	Ioana Baldini, Dennis Wei, Karthikeyan Natesan Ramamurthy, Moninder Singh, Mikhail Yurochkin
<i>Zero-Shot Dense Retrieval with Momentum Adversarial Domain Invariant Representations</i>	Ji Xin, Chenyan Xiong, Ashwin Srinivasan, Ankita Sharma, Damien Jose, Paul N. Bennett
<i>Zero-shot Learning for Grapheme to Phoneme Conversion with Language Ensemble</i>	Xinjian Li, Florian Metze, David R Mortensen, Shinji Watanabe, Alan Black
<i>ZiNet: Linking Chinese Characters Spanning Three Thousand Years</i>	Yang Chi, Fausto Giunchiglia, daqian Shi, Xiaolei Diao, Chuntao Li, Hao Xu

Short Papers

Title	Authors
A Feasibility Study of Answer-Unaware Question Generation for Education	Liam Dugan, Eleni Miltsakaki, Shriyash Kaustubh Upadhyay, Etan Ginsberg, Hannah Gonzalez, DaHyeon Choi, Chuning Yuan, Chris Callison-Burch

A Simple yet Effective Relation Information Guided Approach for Few-Shot Relation Extraction	Yang Liu, Jinpeng Hu, Xiang Wan, Tsung-Hui Chang
Addressing Resource and Privacy Constraints in Semantic Parsing Through Data Augmentation	Kevin Yang, Olivia Deng, Charles Chen, Richard Shin, Subhro Roy, Benjamin Van Durme
An Isotropy Analysis in the Multilingual BERT Embedding Space	Sara Rajaei, Mohammad Taher Pilehvar
Attention Mechanism with Energy-Friendly Operations	Yu Wan, Baosong Yang, Dayiheng Liu, Rong Xiao, Derek F. Wong, Haibo Zhang, Boxing Chen, Lidia S. Chao
Automatic Speech Recognition and Query By Example for Creole Languages Documentation	Cécile Macaire, Didier Schwab, Benjamin Lecouteux, Emmanuel Schang
Auxiliary tasks to boost Biaffine Semantic Dependency Parsing	Marie Candito
Benchmarking Answer Verification Methods for Question Answering-Based Summarization Evaluation Metrics	Daniel Deutsch, Dan Roth
BPE vs. Morphological Segmentation: A Case Study on Machine Translation of Four Polysynthetic Languages	Manuel Mager, Arturo Oncevay, Elisabeth Mager, Katharina Kann, Thang Vu
Combining (Second-Order) Graph-Based and Headed-Span-Based Projective Dependency Parsing	Songlin Yang, Kewei Tu
Combining Static and Contextualised Multilingual Embeddings	Katharina Hämmerl, Jindřich Libovický, Alexander Fraser
Compressing Sentence Representation for Semantic Retrieval via Homomorphic Projective Distillation	Xuandong Zhao, Zhiguo Yu, Ming Wu, Lei Li
Debiasing Event Understanding for Visual Commonsense Tasks	Minji Seo, YeonJoon Jung, Seungtaek Choi, seung-won hwang, Bei Liu
Does BERT really agree ? Fine-grained Analysis of Lexical Dependence on a Syntactic Task	Karim Lasri, Alessandro Lenci, Thierry Poibeau
Domain Generalisation of NMT: Fusing Adapters with Leave-One-Domain-Out Training	Thuy-Trang Vu, Shahram Khadivi, Dinh Phung, Gholamreza Haffari
Dual Context-Guided Continuous Prompt Tuning for Few-Shot Learning	Jie Zhou, Le Tian, Houjin Yu, Zhou Xiao, Hui Su, Jie Zhou
Dynamically Refined Regularization for Improving Cross-corpora Hate Speech Detection	Tulika Bose, Nikolaos Aletras, Irina Illina, Dominique Fohr
EICO: Improving Few-Shot Text Classification via Explicit and Implicit Consistency Regularization	Lei Zhao, cheng yao
Factual Consistency of Multilingual Pretrained Language Models	Constanza Fierro, Anders Søgaard
First the Worst: Finding Better Gender Translations During Beam Search	Danielle Saunders, Rosie Sallis, Bill Byrne
Hierarchical Inductive Transfer for Continual Dialogue Learning	Shaoxiong Feng, Xuancheng Ren, Kan Li, Xu Sun
Improving Zero-Shot Cross-lingual Transfer Between Closely Related Languages by Injecting Character-Level Noise	Noëmi Aeppli, Rico Sennrich

Incremental Intent Detection for Medical Domain with Contrast Replay Networks	Guirong Bai, Shizhu He, Kang Liu, Jun Zhao
Investigating Data Variance in Evaluations of Automatic Machine Translation Metrics	Jiannan Xiang, Huayang Li, Yahui Liu, Lemao Liu, Guoping Huang, Defu Lian, Shuming Shi
Investigating Selective Prediction Approaches Across Several Tasks in IID, OOD, and Adversarial Settings	Neeraj Varshney, Swaroop Mishra, Chitta Baral
Is Whole Word Masking Always Better for Chinese BERT?: Probing on Chinese Grammatical Error Correction	Yong Dai, Linyang Li, Cong Zhou, Zhangyin Feng, Enbo Zhao, Xipeng Qiu, Piji Li, Duyu Tang
Leveraging Expert Guided Adversarial Augmentation For Improving Generalization in Named Entity Recognition	Aaron Reich, Jiaao Chen, Aastha Agrawal, Yanzhe Zhang, Diyi Yang
Mitigating Contradictions in Dialogue Based on Contrastive Learning	Weizhao Li, Junsheng Kong, Ben Liao, Yi Cai
Modular and Parameter-Efficient Multimodal Fusion with Prompting	Sheng Liang, Mengjie Zhao, Hinrich Schuetze
MR-P: A Parallel Decoding Algorithm for Iterative Refinement Non-Autoregressive Translation	Hao Cheng, Zhihua Zhang
MTRec: Multi-Task Learning over BERT for News Recommendation	Qiwei Bi, Jian Li, Lifeng Shang, Xin Jiang, Qun Liu, Hanfang Yang
On Length Divergence Bias in Textual Matching Models	Lan Jiang, Tianshu Lyu, Yankai Lin, Meng Chong, Xiaoyong Lyu, Dawei Yin
Pre-Trained Multilingual Sequence-to-Sequence Models: A Hope for Low-Resource Language Translation?	En-Shiun Annie Lee, Sarubi Thillainathan, Shravan Nayak, Surangika Ranathunga, David Ifeoluwa Adelani, Ruisi Su, Arya D. McCarthy
Prompt Tuning for Discriminative Pre-trained Language Models	Yuan Yao, Bowen Dong, Ao Zhang, Zhengyan Zhang, Ruobing Xie, Zhiyuan Liu, Leyu Lin, Maosong Sun, Jianyong Wang
Read Top News First: A Document Reordering Approach for Multi-Document News Summarization	Chao Zhao, Tenghao Huang, Somnath Basu Roy Chowdhury, Muthu Kumar Chandrasekaran, Kathleen McKeown, Snigdha Chaturvedi
Relevant CommonSense Subgraphs for What if... Procedural Reasoning	Chen Zheng, Parisa Kordjamshidi
Revisiting the Effects of Leakage on Dependency Parsing	Nathaniel Krasner, Miriam Wanner, Antonios Anastasopoulos
Revisiting Uncertainty-based Query Strategies for Active Learning with Transformers	Christopher Schröder, Andreas Niekler, Martin Potthast
Richer Countries and Richer Representations	Kaitlyn Zhou, Kawin Ethayarajh, Dan Jurafsky
RuCCoN: Clinical Concept Normalization in Russian	Alexandr Nesterov, Galina Zubkova, Zulfat Miftahutdinov, Vladimir Kokh, Elena Tutubalina, Artem Shelmanov, Anton M. Alekseev, Manvel Avetisian, Andrey Chertok, Sergey Nikolenko
S ² SQL: Injecting Syntax to Question-Schema Interaction Graph Encoder for Text-to-SQL Parsers	Binyuan Hui, Ruiying Geng, Lihan Wang, Bowen Qin, Yanyang Li, Bowen Li, Jian SUN, Yongbin Li

Selecting Stickers in Open-Domain Dialogue through Multitask Learning	Zhexin Zhang, Yeshuang Zhu, Zhengcong Fei, Jinchao Zhang, Jie Zhou
Single Model Ensemble for Subword Regularized Models in Low-Resource Machine Translation	Sho Takase, Tatsuya Hiraoka, Naoaki Okazaki
Striking a Balance: Alleviating Inconsistency in Pre-trained Models for Symmetric Classification Tasks	Ashutosh Kumar, Aditya Joshi
SyMCoM - Syntactic Measure of Code Mixing A Study Of English-Hindi Code-Mixing	Prashant Kodali, Anmol Goel, Monojit Choudhury, Manish Shrivastava, Ponnurangam Kumaraguru
Towards Responsible Natural Language Annotation for the Varieties of Arabic	A. Stevie Bergman, Mona T. Diab
Transfer Learning and Prediction Consistency for Detecting Offensive Spans of Text	Amir Pouran Ben Veyseh, Ning Xu, Quan Hung Tran, Varun Manjunatha, Franck Dernoncourt, Thien Huu Nguyen
Two Birds with One Stone: Unified Model Learning for Both Recall and Ranking in News Recommendation	Chuhan Wu, Fangzhao Wu, Tao Qi, Yongfeng Huang
Two-Step Question Retrieval for Open-Domain QA	Yeon Seonwoo, Juhee Son, Jiho Jin, Sang-Woo Lee, Ji-Hoon Kim, Jung-Woo Ha, Alice Oh
uFACT: Unfaithful Alien-Corpora Training for Semantically Consistent Data-to-Text Generation	Tisha Anders, Alexandru Coca, Bill Byrne
Unsupervised Chinese Word Segmentation with BERT Oriented Probing and Transformation	Wei Li, Yuhan Song, Qi Su, Yanqiu Shao
Unsupervised Preference-Aware Language Identification	Xingzhang Ren, Baosong Yang, Dayiheng Liu, Haibo Zhang, Xiaoyu Lv, Liang Yao, jun xie
What is wrong with you?: Leveraging User Sentiment for Automatic Dialog Evaluation	Sarik Ghazarian, Behnam Hedayatnia, Alexandros Papangelis, Yang Liu, Dilek Hakkani-Tur
Word-level Perturbation Considering Word Length and Compositional Subwords	Tatsuya Hiraoka, Sho Takase, Kei Uchiumi, Atsushi Keyaki, Naoaki Okazaki