

- 1, TITLE: Investigating label suggestions for opinion mining in German Covid-19 social media  
<https://aclanthology.org/2021.acl-long.1>  
AUTHORS: Tilman Beck, Ji-Ung Lee, Christina Viehmann, Marcus Maurer, Oliver Quiring, Iryna Gurevych  
HIGHLIGHT: This work investigates the use of interactively updated label suggestions to improve upon the efficiency of gathering annotations on the task of opinion mining in German Covid-19 social media data.
- 2, TITLE: How Did This Get Funded?! Automatically Identifying Quirky Scientific Achievements  
<https://aclanthology.org/2021.acl-long.2>  
AUTHORS: Chen Shani, Nadav Borenstein, Dafna Shahaf  
HIGHLIGHT: In this work, we introduce a novel setting in humor mining: automatically detecting funny and unusual scientific papers.
- 3, TITLE: Engage the Public: Poll Question Generation for Social Media Posts  
<https://aclanthology.org/2021.acl-long.3>  
AUTHORS: Zexin Lu, Keyang Ding, Yuji Zhang, Jing Li, Baolin Peng, Lemao Liu  
HIGHLIGHT: This paper presents a novel task to generate poll questions for social media posts.
- 4, TITLE: HateCheck: Functional Tests for Hate Speech Detection Models  
<https://aclanthology.org/2021.acl-long.4>  
AUTHORS: Paul R?ttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, Janet Pierrehumbert  
HIGHLIGHT: To enable more targeted diagnostic insights, we introduce HateCheck, a suite of functional tests for hate speech detection models.
- 5, TITLE: Unified Dual-view Cognitive Model for Interpretable Claim Verification  
<https://aclanthology.org/2021.acl-long.5>  
AUTHORS: Lianwei Wu, Yuan Rao, Yuqian Lan, Ling Sun, Zhaoyin Qi  
HIGHLIGHT: In this paper, we propose a Dual-view model based on the views of Collective and Individual Cognition (CICD) for interpretable claim verification.
- 6, TITLE: DeepRapper: Neural Rap Generation with Rhyme and Rhythm Modeling  
<https://aclanthology.org/2021.acl-long.6>  
AUTHORS: Lanqing Xue, Kaitao Song, Duocai Wu, Xu Tan, Nevin L. Zhang, Tao Qin, Wei-Qiang Zhang, Tie-Yan Liu  
HIGHLIGHT: In this paper, we develop DeepRapper, a Transformer-based rap generation system that can model both rhymes and rhythms.
- 7, TITLE: PENS: A Dataset and Generic Framework for Personalized News Headline Generation  
<https://aclanthology.org/2021.acl-long.7>  
AUTHORS: Xiang Ao, Xiting Wang, Ling Luo, Ying Qiao, Qing He, Xing Xie  
HIGHLIGHT: In this paper, we formulate the personalized news headline generation problem whose goal is to output a user-specific title based on both a user's reading interests and a candidate news body to be exposed to her.
- 8, TITLE: Enhancing Content Preservation in Text Style Transfer Using Reverse Attention and Conditional Layer Normalization  
<https://aclanthology.org/2021.acl-long.8>  
AUTHORS: Dongkyu Lee, Zhiliang Tian, Lanqing Xue, Nevin L. Zhang  
HIGHLIGHT: In this paper, we propose to enhance content preservation by implicitly removing the style information of each token with reverse attention, and thereby retain the content.
- 9, TITLE: Mention Flags (MF): Constraining Transformer-based Text Generators  
<https://aclanthology.org/2021.acl-long.9>  
AUTHORS: Yufei Wang, Ian Wood, Stephen Wan, Mark Dras, Mark Johnson  
HIGHLIGHT: In this paper, we propose Mention Flags (MF), which traces whether lexical constraints are satisfied in the generated outputs in an S2S decoder.
- 10, TITLE: Generalising Multilingual Concept-to-Text NLG with Language Agnostic Delexicalisation  
<https://aclanthology.org/2021.acl-long.10>  
AUTHORS: Giulio Zhou, Gerasimos Lampouras  
HIGHLIGHT: In this paper, we explore the application of multilingual models in concept-to-text and propose Language Agnostic Delexicalisation, a novel delexicalisation method that uses multilingual pretrained embeddings, and employs a character-level post-editing model to inflect words in their correct form during relexicalisation.

- 11, TITLE: Conversations Are Not Flat: Modeling the Dynamic Information Flow across Dialogue Utterances  
<https://aclanthology.org/2021.acl-long.11>  
AUTHORS: Zekang Li, Jinchao Zhang, Zhengcong Fei, Yang Feng, Jie Zhou  
HIGHLIGHT: In this work, we propose the DialoFlow model, in which we introduce a dynamic flow mechanism to model the context flow, and design three training objectives to capture the information dynamics across dialogue utterances by addressing the semantic influence brought about by each utterance in large-scale pre-training.
- 12, TITLE: Dual Slot Selector via Local Reliability Verification for Dialogue State Tracking  
<https://aclanthology.org/2021.acl-long.12>  
AUTHORS: Jinyu Guo, Kai Shuang, Jijie Li, Zihan Wang  
HIGHLIGHT: To address this problem, we devise the two-stage DSS-DST which consists of the Dual Slot Selector based on the current turn dialogue, and the Slot Value Generator based on the dialogue history.
- 13, TITLE: Transferable Dialogue Systems and User Simulators  
<https://aclanthology.org/2021.acl-long.13>  
AUTHORS: Bo-Hsiang Tseng, Yinpei Dai, Florian Kreyszig, Bill Byrne  
HIGHLIGHT: Our goal is to develop a modelling framework that can incorporate new dialogue scenarios through self-play between the two agents.
- 14, TITLE: BoB: BERT Over BERT for Training Persona-based Dialogue Models from Limited Personalized Data  
<https://aclanthology.org/2021.acl-long.14>  
AUTHORS: Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, Ting Liu  
HIGHLIGHT: This work shows how this challenge can be addressed by disentangling persona-based dialogue generation into two sub-tasks with a novel BERT-over-BERT (BoB) model.
- 15, TITLE: GL-GIN: Fast and Accurate Non-Autoregressive Model for Joint Multiple Intent Detection and Slot Filling  
<https://aclanthology.org/2021.acl-long.15>  
AUTHORS: Libo Qin, Fuxuan Wei, Tianbao Xie, Xiao Xu, Wanxiang Che, Ting Liu  
HIGHLIGHT: In this paper, we explore a non-autoregressive model for joint multiple intent detection and slot filling, achieving more fast and accurate.
- 16, TITLE: Accelerating BERT Inference for Sequence Labeling via Early-Exit  
<https://aclanthology.org/2021.acl-long.16>  
AUTHORS: Xiaonan Li, Yunfan Shao, Tianxiang Sun, Hang Yan, Xipeng Qiu, Xuanjing Huang  
HIGHLIGHT: The token-level early-exit brings the gap between training and inference, so we introduce an extra self-sampling fine-tuning stage to alleviate it.
- 17, TITLE: Modularized Interaction Network for Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.17>  
AUTHORS: Fei Li, Zheng Wang, Siu Cheung Hui, Lejian Liao, Dandan Song, Jing Xu, Guoxiu He, Meihuizi Jia  
HIGHLIGHT: In this paper, we propose a novel Modularized Interaction Network (MIN) model which utilizes both segment-level information and word-level dependencies, and incorporates an interaction mechanism to support information sharing between boundary detection and type prediction to enhance the performance for the NER task.
- 18, TITLE: Capturing Event Argument Interaction via A Bi-Directional Entity-Level Recurrent Decoder  
<https://aclanthology.org/2021.acl-long.18>  
AUTHORS: Xi Xiangyu, Wei Ye, Shikun Zhang, Quanxiu Wang, Huixing Jiang, Wei Wu  
HIGHLIGHT: To tackle the above two bottlenecks, we formalize EAE as a Seq2Seq-like learning problem for the first time, where a sentence with a specific event trigger is mapped to a sequence of event argument roles.
- 19, TITLE: UniRE: A Unified Label Space for Entity Relation Extraction  
<https://aclanthology.org/2021.acl-long.19>  
AUTHORS: Yijun Wang, Changzhi Sun, Yuanbin Wu, Hao Zhou, Lei Li, Junchi Yan  
HIGHLIGHT: In this work, we propose to eliminate the different treatment on the two sub-tasks' label spaces.
- 20, TITLE: Refining Sample Embeddings with Relation Prototypes to Enhance Continual Relation Extraction  
<https://aclanthology.org/2021.acl-long.20>  
AUTHORS: Li Cui, Deqing Yang, Jiaxin Yu, Chengwei Hu, Jiayang Cheng, Jingjie Yi, Yanghua Xiao  
HIGHLIGHT: To fully utilize memorized samples, in this paper, we employ relation prototype to extract useful information of each relation.

- 21, TITLE: Contrastive Learning for Many-to-many Multilingual Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.21>  
AUTHORS: Xiao Pan, Mingxuan Wang, Liwei Wu, Lei Li  
HIGHLIGHT: In this work, we aim to build a many-to-many translation system with an emphasis on the quality of non-English language directions.
- 22, TITLE: Understanding the Properties of Minimum Bayes Risk Decoding in Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.22>  
AUTHORS: Mathias Müller, Rico Sennrich  
HIGHLIGHT: In this paper, we empirically investigate the properties of MBR decoding on a number of previously reported biases and failure cases of beam search.
- 23, TITLE: Multi-Head Highly Parallelized LSTM Decoder for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.23>  
AUTHORS: Hongfei Xu, Qiuhui Liu, Josef van Genabith, Deyi Xiong, Meng Zhang  
HIGHLIGHT: To enable sequence-level parallelization of LSTMs, we approximate full LSTM context modelling by computing hidden states and gates with the current input and a simple bag-of-words representation of the preceding tokens context.
- 24, TITLE: A Bidirectional Transformer Based Alignment Model for Unsupervised Word Alignment  
<https://aclanthology.org/2021.acl-long.24>  
AUTHORS: Jingyi Zhang, Josef van Genabith  
HIGHLIGHT: This paper presents a bidirectional Transformer based alignment (BTBA) model for unsupervised learning of the word alignment task.
- 25, TITLE: Learning Language Specific Sub-network for Multilingual Machine Translation  
<https://aclanthology.org/2021.acl-long.25>  
AUTHORS: Zehui Lin, Liwei Wu, Mingxuan Wang, Lei Li  
HIGHLIGHT: In this paper, we propose LaSS to jointly train a single unified multilingual MT model.
- 26, TITLE: Exploring the Efficacy of Automatically Generated Counterfactuals for Sentiment Analysis  
<https://aclanthology.org/2021.acl-long.26>  
AUTHORS: Linyi Yang, Jiazheng Li, Pdraig Cunningham, Yue Zhang, Barry Smyth, Ruihai Dong  
HIGHLIGHT: In this work, we propose an alternative by describing and evaluating an approach to automatically generating counterfactual data for the purpose of data augmentation and explanation.
- 27, TITLE: Bridge-Based Active Domain Adaptation for Aspect Term Extraction  
<https://aclanthology.org/2021.acl-long.27>  
AUTHORS: Zhuang Chen, Tiejun Qian  
HIGHLIGHT: In this paper, we propose a novel active domain adaptation method.
- 28, TITLE: Multimodal Sentiment Detection Based on Multi-channel Graph Neural Networks  
<https://aclanthology.org/2021.acl-long.28>  
AUTHORS: Xiaocui Yang, Shi Feng, Yifei Zhang, Daling Wang  
HIGHLIGHT: In this paper, we propose Multi-channel Graph Neural Networks with Sentiment-awareness (MGNNs) for image-text sentiment detection.
- 29, TITLE: Aspect-Category-Opinion-Sentiment Quadruple Extraction with Implicit Aspects and Opinions  
<https://aclanthology.org/2021.acl-long.29>  
AUTHORS: Hongjie Cai, Rui Xia, Jianfei Yu  
HIGHLIGHT: In this work, we introduce a new task, named Aspect-Category-Opinion-Sentiment (ACOS) Quadruple Extraction, with the goal to extract all aspect-category-opinion-sentiment quadruples in a review sentence and provide full support for aspect-based sentiment analysis with implicit aspects and opinions.
- 30, TITLE: PASS: Perturb-and-Select Summarizer for Product Reviews  
<https://aclanthology.org/2021.acl-long.30>  
AUTHORS: Nadav Oved, Ran Levy  
HIGHLIGHT: We propose the PASS system (Perturb-and-Select Summarizer) that employs a large pre-trained Transformer-based model (T5 in our case), which follows a few-shot fine-tuning scheme.

- 31, TITLE: Deep Differential Amplifier for Extractive Summarization  
<https://aclanthology.org/2021.acl-long.31>  
AUTHORS: Ruipeng Jia, Yanan Cao, Fang Fang, Yuchen Zhou, Zheng Fang, Yanbing Liu, Shi Wang  
HIGHLIGHT: In this paper, we conceptualize the single-document extractive summarization as a rebalance problem and present a deep differential amplifier framework.
- 32, TITLE: Multi-TimeLine Summarization (MTLS): Improving Timeline Summarization by Generating Multiple Summaries  
<https://aclanthology.org/2021.acl-long.32>  
AUTHORS: Yi Yu, Adam Jatowt, Antoine Doucet, Kazunari Sugiyama, Masatoshi Yoshikawa  
HIGHLIGHT: In this paper, we address a novel task, Multiple TimeLine Summarization (MTLS), which extends the flexibility and versatility of Time-Line Summarization (TLS).
- 33, TITLE: Self-Supervised Multimodal Opinion Summarization  
<https://aclanthology.org/2021.acl-long.33>  
AUTHORS: Jinbae Im, Moonki Kim, Hoyeop Lee, Hyunsouk Cho, Sehee Chung  
HIGHLIGHT: To use the abundant information contained in non-text data, we propose a self-supervised multimodal opinion summarization framework called MultimodalSum.
- 34, TITLE: A Training-free and Reference-free Summarization Evaluation Metric via Centrality-weighted Relevance and Self-referenced Redundancy  
<https://aclanthology.org/2021.acl-long.34>  
AUTHORS: Wang Chen, Piji Li, Irwin King  
HIGHLIGHT: To avoid these limitations, we propose a training-free and reference-free summarization evaluation metric.
- 35, TITLE: DESCGEN: A Distantly Supervised Dataset for Generating Entity Descriptions  
<https://aclanthology.org/2021.acl-long.35>  
AUTHORS: Weijia Shi, Mandar Joshi, Luke Zettlemoyer  
HIGHLIGHT: We introduce DESCGEN: given mentions spread over multiple documents, the goal is to generate an entity summary description.
- 36, TITLE: Introducing Orthogonal Constraint in Structural Probes  
<https://aclanthology.org/2021.acl-long.36>  
AUTHORS: Tomasz Limisiewicz, David Marecek  
HIGHLIGHT: In this work, we introduce a new type of structural probing, where the linear projection is decomposed into 1. iso-morphic space rotation; 2. linear scaling that identifies and scales the most relevant dimensions.
- 37, TITLE: Hidden Killer: Invisible Textual Backdoor Attacks with Syntactic Trigger  
<https://aclanthology.org/2021.acl-long.37>  
AUTHORS: Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, Maosong Sun  
HIGHLIGHT: In this paper, we propose to use the syntactic structure as the trigger in textual backdoor attacks.
- 38, TITLE: Examining the Inductive Bias of Neural Language Models with Artificial Languages  
<https://aclanthology.org/2021.acl-long.38>  
AUTHORS: Jennifer C. White, Ryan Cotterell  
HIGHLIGHT: We propose a novel method for investigating the inductive biases of language models using artificial languages.
- 39, TITLE: Explaining Contextualization in Language Models using Visual Analytics  
<https://aclanthology.org/2021.acl-long.39>  
AUTHORS: Rita Sevastjanova, Aikaterini-Lida Kalouli, Christin Beck, Hanna Schfer, Mennatallah El-Assady  
HIGHLIGHT: In this paper, we contribute to the current efforts of explaining such models by exploring the continuum between function and content words with respect to contextualization in BERT, based on linguistically-informed insights.
- 40, TITLE: Improving the Faithfulness of Attention-based Explanations with Task-specific Information for Text Classification  
<https://aclanthology.org/2021.acl-long.40>  
AUTHORS: George Chrysostomou, Nikolaos Aletras  
HIGHLIGHT: In this paper, we seek to improve the faithfulness of attention-based explanations for text classification.
- 41, TITLE: Generating Landmark Navigation Instructions from Maps as a Graph-to-Text Problem

- <https://aclanthology.org/2021.acl-long.41>  
AUTHORS: Raphael Schumann, Stefan Riezler  
HIGHLIGHT: We present a neural model that takes OpenStreetMap representations as input and learns to generate navigation instructions that contain visible and salient landmarks from human natural language instructions.
- 42, TITLE: E2E-VLP: End-to-End Vision-Language Pre-training Enhanced by Visual Learning  
<https://aclanthology.org/2021.acl-long.42>  
AUTHORS: Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, Fei Huang  
HIGHLIGHT: In this paper, we propose the first end-to-end vision-language pre-trained model for both V+L understanding and generation, namely E2E-VLP, where we build a unified Transformer framework to jointly learn visual representation, and semantic alignments between image and text.
- 43, TITLE: Learning Relation Alignment for Calibrated Cross-modal Retrieval  
<https://aclanthology.org/2021.acl-long.43>  
AUTHORS: Shuhuai Ren, Junyang Lin, Guangxiang Zhao, Rui Men, An Yang, Jingren Zhou, Xu Sun, Hongxia Yang  
HIGHLIGHT: In response, we present Inter-modal Alignment on Intra-modal Self-attentions (IAIS), a regularized training method to optimize the ISD and calibrate intra-modal self-attentions from the two modalities mutually via inter-modal alignment.
- 44, TITLE: KM-BART: Knowledge Enhanced Multimodal BART for Visual Commonsense Generation  
<https://aclanthology.org/2021.acl-long.44>  
AUTHORS: Yiran Xing, Zai Shi, Zhao Meng, Gerhard Lakemeyer, Yunpu Ma, Roger Wattenhofer  
HIGHLIGHT: We present Knowledge Enhanced Multimodal BART (KM-BART), which is a Transformer-based sequence-to-sequence model capable of reasoning about commonsense knowledge from multimodal inputs of images and texts.
- 45, TITLE: Cascaded Head-colliding Attention  
<https://aclanthology.org/2021.acl-long.45>  
AUTHORS: Lin Zheng, Zhiyong Wu, Lingpeng Kong  
HIGHLIGHT: We present cascaded head-colliding attention (CODA) which explicitly models the interactions between attention heads through a hierarchical variational distribution.
- 46, TITLE: Structural Knowledge Distillation: Tractably Distilling Information for Structured Predictor  
<https://aclanthology.org/2021.acl-long.46>  
AUTHORS: Xinyu Wang, Yong Jiang, Zhaohui Yan, Zixia Jia, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, Kewei Tu  
HIGHLIGHT: In this paper, we derive a factorized form of the knowledge distillation objective for structured prediction, which is tractable for many typical choices of the teacher and student models.
- 47, TITLE: Parameter-efficient Multi-task Fine-tuning for Transformers via Shared Hypernetworks  
<https://aclanthology.org/2021.acl-long.47>  
AUTHORS: Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, James Henderson  
HIGHLIGHT: In this paper, we show that we can learn adapter parameters for all layers and tasks by generating them using shared hypernetworks, which condition on task, adapter position, and layer id in a transformer model.
- 48, TITLE: COSY: COunterfactual SYntax for Cross-Lingual Understanding  
<https://aclanthology.org/2021.acl-long.48>  
AUTHORS: Sicheng Yu, Hao Zhang, Yulei Niu, Qianru Sun, Jing Jiang  
HIGHLIGHT: We tackle this issue by incorporating language-agnostic information, specifically, universal syntax such as dependency relations and POS tags, into language models, based on the observation that universal syntax is transferable across different languages.
- 49, TITLE: OoMMix: Out-of-manifold Regularization in Contextual Embedding Space for Text Classification  
<https://aclanthology.org/2021.acl-long.49>  
AUTHORS: Seonghyeon Lee, Dongha Lee, Hwanjo Yu  
HIGHLIGHT: In this work, we propose a new approach, called OoMMix, to finding and regularizing the remainder of the space, referred to as out-of-manifold, which cannot be accessed through the words.
- 50, TITLE: Understanding and Countering Stereotypes: A Computational Approach to the Stereotype Content Model  
<https://aclanthology.org/2021.acl-long.50>  
AUTHORS: Kathleen C. Fraser, Isar Nejadgholi, Svetlana Kiritchenko  
HIGHLIGHT: In this work, we present a computational approach to interpreting stereotypes in text through the Stereotype Content Model (SCM), a comprehensive causal theory from social psychology.

- 51, TITLE: Structurizing Misinformation Stories via Rationalizing Fact-Checks  
<https://aclanthology.org/2021.acl-long.51>  
AUTHORS: Shan Jiang, Christo Wilson  
HIGHLIGHT: This paper aims to structurize these misinformation stories by leveraging fact-check articles.
- 52, TITLE: Modeling Language Usage and Listener Engagement in Podcasts  
<https://aclanthology.org/2021.acl-long.52>  
AUTHORS: Sravana Reddy, Mariya Lazarova, Yongze Yu, Rosie Jones  
HIGHLIGHT: In this paper, we investigate how various factors - vocabulary diversity, distinctiveness, emotion, and syntax, among others - correlate with engagement, based on analysis of the creators' written descriptions and transcripts of the audio.
- 53, TITLE: Breaking Down the Invisible Wall of Informal Fallacies in Online Discussions  
<https://aclanthology.org/2021.acl-long.53>  
AUTHORS: Saumya Sahai, Oana Balalau, Roxana Horincar  
HIGHLIGHT: In this paper, we study the most frequent fallacies on Reddit, and we present them using the pragma-dialectical theory of argumentation.
- 54, TITLE: SocAoG: Incremental Graph Parsing for Social Relation Inference in Dialogues  
<https://aclanthology.org/2021.acl-long.54>  
AUTHORS: Liang Qiu, Yuan Liang, Yizhou Zhao, Pan Lu, Baolin Peng, Zhou Yu, Ying Nian Wu, Song-Chun Zhu  
HIGHLIGHT: We model the social network as an And-or Graph, named SocAoG, for the consistency of relations among a group and leveraging attributes as inference cues.
- 55, TITLE: TicketTalk: Toward human-level performance with end-to-end, transaction-based dialog systems  
<https://aclanthology.org/2021.acl-long.55>  
AUTHORS: Bill Byrne, Karthik Krishnamoorthi, Saravanan Ganesh, Mihir Kale  
HIGHLIGHT: We present a data-driven, end-to-end approach to transaction-based dialog systems that performs at near-human levels in terms of verbal response quality and factual grounding accuracy.
- 56, TITLE: Improving Dialog Systems for Negotiation with Personality Modeling  
<https://aclanthology.org/2021.acl-long.56>  
AUTHORS: Runzhe Yang, Jingxiao Chen, Karthik Narasimhan  
HIGHLIGHT: In this paper, we explore the ability to model and infer personality types of opponents, predict their responses, and use this information to adapt a dialog agent's high-level strategy in negotiation tasks.
- 57, TITLE: Learning from Perturbations: Diverse and Informative Dialogue Generation with Inverse Adversarial Training  
<https://aclanthology.org/2021.acl-long.57>  
AUTHORS: Wangchunshu Zhou, Qifei Li, Chenle Li  
HIGHLIGHT: In this paper, we propose Inverse Adversarial Training (IAT) algorithm for training neural dialogue systems to avoid generic responses and model dialogue history better.
- 58, TITLE: Increasing Faithfulness in Knowledge-Grounded Dialogue with Controllable Features  
<https://aclanthology.org/2021.acl-long.58>  
AUTHORS: Hannah Rashkin, David Reitter, Gaurav Singh Tomar, Dipanjan Das  
HIGHLIGHT: We propose different evaluation measures to disentangle these different styles of responses by quantifying the informativeness and objectivity.
- 59, TITLE: CitationIE: Leveraging the Citation Graph for Scientific Information Extraction  
<https://aclanthology.org/2021.acl-long.59>  
AUTHORS: Vijay Viswanathan, Graham Neubig, Pengfei Liu  
HIGHLIGHT: In contrast to prior work, we augment our text representations by leveraging a complementary source of document context: the citation graph of referential links between citing and cited papers.
- 60, TITLE: From Discourse to Narrative: Knowledge Projection for Event Relation Extraction  
<https://aclanthology.org/2021.acl-long.60>  
AUTHORS: Jialong Tang, Hongyu Lin, Meng Liao, Yaojie Lu, Xianpei Han, Le Sun, Weijian Xie, Jin Xu  
HIGHLIGHT: In this paper, we propose a knowledge projection paradigm for event relation extraction: projecting discourse knowledge to narratives by exploiting the commonalities between them.

61, TITLE: AdvPicker: Effectively Leveraging Unlabeled Data via Adversarial Discriminator for Cross-Lingual NER  
<https://aclanthology.org/2021.acl-long.61>  
AUTHORS: Weile Chen, Huiqiang Jiang, Qianhui Wu, Björn Karlsson, Yi Guan  
HIGHLIGHT: We design an adversarial learning framework in which an encoder learns entity domain knowledge from labeled source-language data and better shared features are captured via adversarial training - where a discriminator selects less language-dependent target-language data via similarity to the source language.

62, TITLE: Compare to The Knowledge: Graph Neural Fake News Detection with External Knowledge  
<https://aclanthology.org/2021.acl-long.62>  
AUTHORS: Linmei Hu, Tianchi Yang, Luhao Zhang, Wanjun Zhong, Duyu Tang, Chuan Shi, Nan Duan, Ming Zhou  
HIGHLIGHT: In this paper, we propose a novel end-to-end graph neural model called CompareNet, which compares the news to the knowledge base (KB) through entities for fake news detection.

63, TITLE: Discontinuous Named Entity Recognition as Maximal Clique Discovery  
<https://aclanthology.org/2021.acl-long.63>  
AUTHORS: Yucheng Wang, Bowen Yu, Hongsong Zhu, Tingwen Liu, Nan Yu, Limin Sun  
HIGHLIGHT: To solve this problem, we first construct a segment graph for each sentence, in which each node denotes a segment (a continuous entity on its own, or a part of discontinuous entities), and an edge links two nodes that belong to the same entity.

64, TITLE: LNN-EL: A Neuro-Symbolic Approach to Short-text Entity Linking  
<https://aclanthology.org/2021.acl-long.64>  
AUTHORS: Hang Jiang, Sairam Gurajada, Qiuhaio Lu, Sumit Neelam, Lucian Popa, Prithviraj Sen, Yunyao Li, Alexander Gray  
HIGHLIGHT: Here, we take a different, neuro-symbolic approach that combines the advantages of using interpretable rules based on first-order logic with the performance of neural learning.

65, TITLE: Do Context-Aware Translation Models Pay the Right Attention?  
<https://aclanthology.org/2021.acl-long.65>  
AUTHORS: Kayo Yin, Patrick Fernandes, Danish Pruthi, Aditi Chaudhary, André F. T. Martins, Graham Neubig  
HIGHLIGHT: To answer these questions, we introduce SCAT (Supporting Context for Ambiguous Translations), a new English-French dataset comprising supporting context words for 14K translations that professional translators found useful for pronoun disambiguation.

66, TITLE: Adapting High-resource NMT Models to Translate Low-resource Related Languages without Parallel Data  
<https://aclanthology.org/2021.acl-long.66>  
AUTHORS: Wei-Jen Ko, Ahmed El-Kishky, Adithya Renduchintala, Vishrav Chaudhary, Naman Goyal, Francisco Guzmán, Pascale Fung, Philipp Koehn, Mona Diab  
HIGHLIGHT: In this work, we exploit this linguistic overlap to facilitate translating to and from a low-resource language with only monolingual data, in addition to any parallel data in the related high-resource language.

67, TITLE: Bilingual Lexicon Induction via Unsupervised Bitext Construction and Word Alignment  
<https://aclanthology.org/2021.acl-long.67>  
AUTHORS: Haoyue Shi, Luke Zettlemoyer, Sida I. Wang  
HIGHLIGHT: In this paper, we show it is possible to produce much higher quality lexicons with methods that combine (1) unsupervised bitext mining and (2) unsupervised word alignment.

68, TITLE: Multilingual Speech Translation from Efficient Finetuning of Pretrained Models  
<https://aclanthology.org/2021.acl-long.68>  
AUTHORS: Xian Li, Changhan Wang, Yun Tang, Chau Tran, Yuqing Tang, Juan Pino, Alexei Baevski, Alexis Conneau, Michael Auli  
HIGHLIGHT: We present a simple yet effective approach to build multilingual speech-to-text (ST) translation through efficient transfer learning from a pretrained speech encoder and text decoder.

69, TITLE: Learning Faithful Representations of Causal Graphs  
<https://aclanthology.org/2021.acl-long.69>  
AUTHORS: Ananth Balashankar, Lakshminarayanan Subramanian  
HIGHLIGHT: In this paper, we define the faithfulness property of contextual embeddings to capture geometric distance-based properties of directed acyclic causal graphs.

- 70, TITLE: What Context Features Can Transformer Language Models Use?  
<https://aclanthology.org/2021.acl-long.70>  
AUTHORS: Joe O'Connor, Jacob Andreas  
HIGHLIGHT: We describe a series of experiments that measure usable information by selectively ablating lexical and structural information in transformer language models trained on English Wikipedia.
- 71, TITLE: Integrated Directional Gradients: Feature Interaction Attribution for Neural NLP Models  
<https://aclanthology.org/2021.acl-long.71>  
AUTHORS: Sandipan Sikdar, Parantapa Bhattacharya, Kieran Heese  
HIGHLIGHT: In this paper, we introduce Integrated Directional Gradients (IDG), a method for attributing importance scores to groups of features, indicating their relevance to the output of a neural network model for a given input.
- 72, TITLE: DeCLUTR: Deep Contrastive Learning for Unsupervised Textual Representations  
<https://aclanthology.org/2021.acl-long.72>  
AUTHORS: John Giorgi, Osvald Nitski, Bo Wang, Gary Bader  
HIGHLIGHT: In this paper, we present DeCLUTR: Deep Contrastive Learning for Unsupervised Textual Representations.
- 73, TITLE: XLPT-AMR: Cross-Lingual Pre-Training via Multi-Task Learning for Zero-Shot AMR Parsing and Text Generation  
<https://aclanthology.org/2021.acl-long.73>  
AUTHORS: Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, Guodong Zhou  
HIGHLIGHT: Upon the availability of English AMR dataset and English-to- X parallel datasets, in this paper we propose a novel cross-lingual pre-training approach via multi-task learning (MTL) for both zeroshot AMR parsing and AMR-to-text generation.
- 74, TITLE: Span-based Semantic Parsing for Compositional Generalization  
<https://aclanthology.org/2021.acl-long.74>  
AUTHORS: Jonathan Herzig, Jonathan Berant  
HIGHLIGHT: In this work, we posit that a span-based parser should lead to better compositional generalization.
- 75, TITLE: Compositional Generalization and Natural Language Variation: Can a Semantic Parsing Approach Handle Both?  
<https://aclanthology.org/2021.acl-long.75>  
AUTHORS: Peter Shaw, Ming-Wei Chang, Panupong Pasupat, Kristina Toutanova  
HIGHLIGHT: In this work we ask: can we develop a semantic parsing approach that handles both natural language variation and compositional generalization?
- 76, TITLE: A Targeted Assessment of Incremental Processing in Neural Language Models and Humans  
<https://aclanthology.org/2021.acl-long.76>  
AUTHORS: Ethan Wilcox, Pranali Vani, Roger Levy  
HIGHLIGHT: We present a targeted, scaled-up comparison of incremental processing in humans and neural language models by collecting by-word reaction time data for sixteen different syntactic test suites across a range of structural phenomena.
- 77, TITLE: The Possible, the Plausible, and the Desirable: Event-Based Modality Detection for Language Processing  
<https://aclanthology.org/2021.acl-long.77>  
AUTHORS: Valentina Pyatkin, Shoval Sadde, Aynat Rubinstein, Paul Portner, Reut Tsarfay  
HIGHLIGHT: This work builds on the theoretical foundations of the Georgetown Gradable Modal Expressions (GME) work by Rubinstein et al. (2013) to propose an event-based modality detection task where modal expressions can be words of any syntactic class and sense labels are drawn from a comprehensive taxonomy which harmonizes the modal concepts contributed by the different studies.
- 78, TITLE: To POS Tag or Not to POS Tag: The Impact of POS Tags on Morphological Learning in Low-Resource Settings  
<https://aclanthology.org/2021.acl-long.78>  
AUTHORS: Sarah Moeller, Ling Liu, Mans Hulden  
HIGHLIGHT: This paper describes an empirical study about the effect that POS tags have on two computational morphological tasks with the Transformer architecture.
- 79, TITLE: Prosodic segmentation for parsing spoken dialogue  
<https://aclanthology.org/2021.acl-long.79>  
AUTHORS: Elizabeth Nielsen, Mark Steedman, Sharon Goldwater



**HIGHLIGHT:** We investigate how prosody affects a parser that receives an entire dialogue turn as input (a turn-based model), instead of gold standard pre-segmented SUs (an SU-based model).

80, **TITLE:** VoxPopuli: A Large-Scale Multilingual Speech Corpus for Representation Learning, Semi-Supervised Learning and Interpretation

<https://aclanthology.org/2021.acl-long.80>

**AUTHORS:** Changhan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, Emmanuel Dupoux

**HIGHLIGHT:** We introduce VoxPopuli, a large-scale multilingual corpus providing 400K hours of unlabeled speech data in 23 languages.

81, **TITLE:** Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets

<https://aclanthology.org/2021.acl-long.81>

**AUTHORS:** Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, Hanna Wallach

**HIGHLIGHT:** We examine four such benchmarks constructed for two NLP tasks: language modeling and coreference resolution.

82, **TITLE:** Robust Knowledge Graph Completion with Stacked Convolutions and a Student Re-Ranking Network

<https://aclanthology.org/2021.acl-long.82>

**AUTHORS:** Justin Lovelace, Denis Newman-Griffis, Shikhar Vashishth, Jill Fain Lehman, Carolyn Ros?

**HIGHLIGHT:** We develop a deep convolutional network that utilizes textual entity representations and demonstrate that our model outperforms recent KG completion methods in this challenging setting.

83, **TITLE:** A DQN-based Approach to Finding Precise Evidences for Fact Verification

<https://aclanthology.org/2021.acl-long.83>

**AUTHORS:** Hai Wan, Haicheng Chen, Jianfeng Du, Weilin Luo, Rongzhen Ye

**HIGHLIGHT:** Inspired by the strong exploration ability of the deep Q-learning network (DQN), we propose a DQN-based approach to retrieval of precise evidences.

84, **TITLE:** The Art of Abstention: Selective Prediction and Error Regularization for Natural Language Processing

<https://aclanthology.org/2021.acl-long.84>

**AUTHORS:** Ji Xin, Raphael Tang, Yaoliang Yu, Jimmy Lin

**HIGHLIGHT:** To fill this void in the literature, we study in this paper selective prediction for NLP, comparing different models and confidence estimators.

85, **TITLE:** Unsupervised Out-of-Domain Detection via Pre-trained Transformers

<https://aclanthology.org/2021.acl-long.85>

**AUTHORS:** Keyang Xu, Tongzheng Ren, Shikun Zhang, Yihao Feng, Caiming Xiong

**HIGHLIGHT:** Our work tackles the problem of detecting out-of-domain samples with only unsupervised in-domain data.

86, **TITLE:** MATE-KD: Masked Adversarial TEXT, a Companion to Knowledge Distillation

<https://aclanthology.org/2021.acl-long.86>

**AUTHORS:** Ahmad Rashid, Vasileios Lioutas, Mehdi Rezagholizadeh

**HIGHLIGHT:** We present MATE-KD, a novel text-based adversarial training algorithm which improves the performance of knowledge distillation.

87, **TITLE:** Selecting Informative Contexts Improves Language Model Fine-tuning

<https://aclanthology.org/2021.acl-long.87>

**AUTHORS:** Richard Antonello, Nicole Beckage, Javier Turek, Alexander Huth

**HIGHLIGHT:** Here we present a general fine-tuning method that we call information gain filtration for improving the overall training efficiency and final performance of language model fine-tuning.

88, **TITLE:** Explainable Prediction of Text Complexity: The Missing Preliminaries for Text Simplification

<https://aclanthology.org/2021.acl-long.88>

**AUTHORS:** Cristina Garbacea, Mengtian Guo, Samuel Carton, Qiaozhu Mei

**HIGHLIGHT:** We show that text simplification can be decomposed into a compact pipeline of tasks to ensure the transparency and explainability of the process.

89, **TITLE:** Multi-Task Retrieval for Knowledge-Intensive Tasks

<https://aclanthology.org/2021.acl-long.89>

AUTHORS: Jean Maillard, Vladimir Karpukhin, Fabio Petroni, Wen-tau Yih, Barlas Oguz, Veselin Stoyanov, Gargi Ghosh  
HIGHLIGHT: Driven by the question of whether a neural retrieval model can be `_universal_` and perform robustly on a wide variety of problems, we propose a multi-task trained model.

90, TITLE: When Do You Need Billions of Words of Pretraining Data?  
<https://aclanthology.org/2021.acl-long.90>  
AUTHORS: Yian Zhang, Alex Warstadt, Xiaocheng Li, Samuel R. Bowman  
HIGHLIGHT: To explore this question, we adopt five styles of evaluation: classifier probing, information-theoretic probing, unsupervised relative acceptability judgments, unsupervised language model knowledge probing, and fine-tuning on NLU tasks.

91, TITLE: Analyzing the Source and Target Contributions to Predictions in Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.91>  
AUTHORS: Elena Voita, Rico Sennrich, Ivan Titov  
HIGHLIGHT: We argue that this relative contribution can be evaluated by adopting a variant of Layerwise Relevance Propagation (LRP).

92, TITLE: Comparing Test Sets with Item Response Theory  
<https://aclanthology.org/2021.acl-long.92>  
AUTHORS: Clara Vania, Phu Mon Htut, William Huang, Dhara Mungra, Richard Yuanzhe Pang, Jason Phang, Haokun Liu, Kyunghyun Cho, Samuel R. Bowman  
HIGHLIGHT: To measure this uniformly across datasets, we draw on Item Response Theory and evaluate 29 datasets using predictions from 18 pretrained Transformer models on individual test examples.

93, TITLE: Uncovering Constraint-Based Behavior in Neural Models via Targeted Fine-Tuning  
<https://aclanthology.org/2021.acl-long.93>  
AUTHORS: Forrest Davis, Marten van Schijndel  
HIGHLIGHT: Our results suggest that models need to learn both the linguistic constraints in a language and their relative ranking, with mismatches in either producing non-human-like behavior.

94, TITLE: More Identifiable yet Equally Performant Transformers for Text Classification  
<https://aclanthology.org/2021.acl-long.94>  
AUTHORS: Rishabh Bhardwaj, Navonil Majumder, Soujanya Poria, Eduard Hovy  
HIGHLIGHT: In this work, we provide deeper theoretical analysis and empirical observations on the identifiability of attention weights.

95, TITLE: AugNLG: Few-shot Natural Language Generation using Self-trained Data Augmentation  
<https://aclanthology.org/2021.acl-long.95>  
AUTHORS: Xinnuo Xu, Guoyin Wang, Young-Bum Kim, Sungjin Lee  
HIGHLIGHT: This paper proposes AugNLG, a novel data augmentation approach that combines a self-trained neural retrieval model with a few-shot learned NLU model, to automatically create MR-to-Text data from open-domain texts.

96, TITLE: Can vectors read minds better than experts? Comparing data augmentation strategies for the automated scoring of children's mindreading ability  
<https://aclanthology.org/2021.acl-long.96>  
AUTHORS: Venelin Kovatchev, Phillip Smith, Mark Lee, Rory Devine  
HIGHLIGHT: In this paper we implement and compare 7 different data augmentation strategies for the task of automatic scoring of children's ability to understand others' thoughts, feelings, and desires (or "mindreading").

97, TITLE: A Dataset and Baselines for Multilingual Reply Suggestion  
<https://aclanthology.org/2021.acl-long.97>  
AUTHORS: Mozhi Zhang, Wei Wang, Budhaditya Deb, Guoqing Zheng, Milad Shokouhi, Ahmed Hassan Awadallah  
HIGHLIGHT: Instead, we present MRS, a multilingual reply suggestion dataset with ten languages.

98, TITLE: What Ingredients Make for an Effective Crowdsourcing Protocol for Difficult NLU Data Collection Tasks?  
<https://aclanthology.org/2021.acl-long.98>  
AUTHORS: Nikita Nangia, Saku Sugawara, Harsh Trivedi, Alex Warstadt, Clara Vania, Samuel R. Bowman  
HIGHLIGHT: In this paper, we compare the efficacy of interventions that have been proposed in prior work as ways of improving data quality.

99, TITLE: Align Voting Behavior with Public Statements for Legislator Representation Learning

<https://aclanthology.org/2021.acl-long.99>

AUTHORS: Xinyi Mou, Zhongyu Wei, Lei Chen, Shangyi Ning, Yancheng He, Changjian Jiang, Xuanjing Huang  
HIGHLIGHT: In order to mitigate these two problems, we explore to incorporate both voting behavior and public statements on Twitter to jointly model legislators.

100, TITLE: Measure and Evaluation of Semantic Divergence across Two Languages

<https://aclanthology.org/2021.acl-long.100>

AUTHORS: Syrielle Montariol, Alexandre Allauzen

HIGHLIGHT: In this paper, we propose to track these divergences by comparing the evolution of a word and its translation across two languages.

101, TITLE: Improving Zero-Shot Translation by Disentangling Positional Information

<https://aclanthology.org/2021.acl-long.101>

AUTHORS: Danni Liu, Jan Niehues, James Cross, Francisco Guzmán, Xian Li

HIGHLIGHT: We demonstrate that a main factor causing the language-specific representations is the positional correspondence to input tokens. We show that this can be easily alleviated by removing residual connections in an encoder layer.

102, TITLE: Common Sense Beyond English: Evaluating and Improving Multilingual Language Models for Commonsense Reasoning

<https://aclanthology.org/2021.acl-long.102>

AUTHORS: Bill Yuchen Lin, Seyeon Lee, Xiaoyang Qiao, Xiang Ren

HIGHLIGHT: We aim to evaluate and improve popular multilingual language models (ML-LMs) to help advance commonsense reasoning (CSR) beyond English.

103, TITLE: Attention Calibration for Transformer in Neural Machine Translation

<https://aclanthology.org/2021.acl-long.103>

AUTHORS: Yu Lu, Jiali Zeng, Jiajun Zhang, Shuangzhi Wu, Mu Li

HIGHLIGHT: In this paper, we propose to calibrate the attention weights by introducing a mask perturbation model that automatically evaluates each input's contribution to the model outputs.

104, TITLE: Diverse Pretrained Context Encodings Improve Document Translation

<https://aclanthology.org/2021.acl-long.104>

AUTHORS: Domenic Donato, Lei Yu, Chris Dyer

HIGHLIGHT: We propose a new architecture for adapting a sentence-level sequence-to-sequence transformer by incorporating multiple pre-trained document context signals and assess the impact on translation performance of (1) different pretraining approaches for generating these signals, (2) the quantity of parallel data for which document context is available, and (3) conditioning on source, target, or source and target contexts.

105, TITLE: Exploiting Language Relatedness for Low Web-Resource Language Model Adaptation: An Indic Languages Study

<https://aclanthology.org/2021.acl-long.105>

AUTHORS: Yash Khemchandani, Sarvesh Mehtani, Vaidehi Patil, Abhijeet Awasthi, Partha Talukdar, Sunita Sarawagi

HIGHLIGHT: In this paper we argue that relatedness among languages in a language family may be exploited to overcome some of the corpora limitations of LRLs, and propose RelateLM.

106, TITLE: On Finding the K-best Non-projective Dependency Trees

<https://aclanthology.org/2021.acl-long.106>

AUTHORS: Ran Zmigrod, Tim Vieira, Ryan Cotterell

HIGHLIGHT: In this paper, we provide a simplification of the K-best spanning tree algorithm of Camerini et al. (1980).

107, TITLE: Towards Argument Mining for Social Good: A Survey

<https://aclanthology.org/2021.acl-long.107>

AUTHORS: Eva Maria Vecchi, Neele Falk, Iman Jundi, Gabriella Lapesa

HIGHLIGHT: We propose a novel definition of argument quality which is integrated with that of deliberative quality from the Social Science literature.

108, TITLE: Automated Generation of Storytelling Vocabulary from Photographs for use in AAC

<https://aclanthology.org/2021.acl-long.108>

AUTHORS: Mauricio Fontana de Vargas, Karyn Moffatt

HIGHLIGHT: We contribute a novel method for generating context-related vocabulary from photographs of personally relevant events aimed at supporting people with language impairments in retelling their past experiences.

109, TITLE: CLIP: A Dataset for Extracting Action Items for Physicians from Hospital Discharge Notes  
<https://aclanthology.org/2021.acl-long.109>  
AUTHORS: James Mullenbach, Yada Pruksachatkun, Sean Adler, Jennifer Seale, Jordan Swartz, Greg McKelvey, Hui Dai, Yi Yang, David Sontag  
HIGHLIGHT: In this work, we describe our creation of a dataset of clinical action items annotated over MIMIC-III, the largest publicly available dataset of real clinical notes.

110, TITLE: Assessing Emoji Use in Modern Text Processing Tools  
<https://aclanthology.org/2021.acl-long.110>  
AUTHORS: Abu Awal Md Shoeb, Gerard de Melo  
HIGHLIGHT: In this study, we assess this support by considering test sets of tweets with emojis, based on which we perform a series of experiments investigating the ability of prominent NLP and text processing tools to adequately process them.

111, TITLE: Select, Extract and Generate: Neural Keyphrase Generation with Layer-wise Coverage Attention  
<https://aclanthology.org/2021.acl-long.111>  
AUTHORS: Wasi Ahmad, Xiao Bai, Soomin Lee, Kai-Wei Chang  
HIGHLIGHT: To overcome this limitation, we propose SEG-Net, a neural keyphrase generation model that is composed of two major components, (1) a selector that selects the salient sentences in a document and (2) an extractor-generator that jointly extracts and generates keyphrases from the selected sentences.

112, TITLE: Factorising Meaning and Form for Intent-Preserving Paraphrasing  
<https://aclanthology.org/2021.acl-long.112>  
AUTHORS: Tom Hosking, Mirella Lapata  
HIGHLIGHT: We propose a method for generating paraphrases of English questions that retain the original intent but use a different surface form.

113, TITLE: AggGen: Ordering and Aggregating while Generating  
<https://aclanthology.org/2021.acl-long.113>  
AUTHORS: Xinnuo Xu, Ondrej Dušek, Verena Rieser, Ioannis Konstas  
HIGHLIGHT: We present AggGen (pronounced 'again') a data-to-text model which re-introduces two explicit sentence planning stages into neural data-to-text systems: input ordering and input aggregation.

114, TITLE: Reflective Decoding: Beyond Unidirectional Generation with Off-the-Shelf Language Models  
<https://aclanthology.org/2021.acl-long.114>  
AUTHORS: Peter West, Ximing Lu, Ari Holtzman, Chandra Bhagavatula, Jena D. Hwang, Yejin Choi  
HIGHLIGHT: In this paper, we present Reflective Decoding, a novel unsupervised algorithm that allows for direct application of unidirectional LMs to non-sequential tasks.

115, TITLE: Towards Table-to-Text Generation with Numerical Reasoning  
<https://aclanthology.org/2021.acl-long.115>  
AUTHORS: Lya Hulliyatus Suadaa, Hidetaka Kamigaito, Kotaro Funakoshi, Manabu Okumura, Hiroya Takamura  
HIGHLIGHT: In this paper, we propose a framework consisting of a pre-trained model and a copy mechanism.

116, TITLE: BACO: A Background Knowledge- and Content-Based Framework for Citing Sentence Generation  
<https://aclanthology.org/2021.acl-long.116>  
AUTHORS: Yubin Ge, Ly Dinh, Xiaofeng Liu, Jinsong Su, Ziyao Lu, Ante Wang, Jana Diesner  
HIGHLIGHT: In this paper, we focus on the problem of citing sentence generation, which entails generating a short text to capture the salient information in a cited paper and the connection between the citing and cited paper.

117, TITLE: Language Model as an Annotator: Exploring DialoGPT for Dialogue Summarization  
<https://aclanthology.org/2021.acl-long.117>  
AUTHORS: Xiachong Feng, Xiaocheng Feng, Libo Qin, Bing Qin, Ting Liu  
HIGHLIGHT: In this paper, we show how DialoGPT, a pre-trained model for conversational response generation, can be developed as an unsupervised dialogue annotator, which takes advantage of dialogue background knowledge encoded in DialoGPT.

118, TITLE: Challenges in Information-Seeking QA: Unanswerable Questions and Paragraph Retrieval  
<https://aclanthology.org/2021.acl-long.118>  
AUTHORS: Akari Asai, Eunsol Choi

**HIGHLIGHT:** Our controlled experiments suggest two headrooms - paragraph selection and answerability prediction, i.e. whether the paired evidence document contains the answer to the query or not.

119, **TITLE:** A Gradually Soft Multi-Task and Data-Augmented Approach to Medical Question Understanding  
<https://aclanthology.org/2021.acl-long.119>  
**AUTHORS:** Khalil Mrini, Franck Dernoncourt, Seunghyun Yoon, Trung Bui, Walter Chang, Emilia Farcas, Ndapa Nakashole  
**HIGHLIGHT:** To alleviate this problem, we propose a novel Multi-Task Learning (MTL) method with data augmentation for medical question understanding.

120, **TITLE:** Leveraging Type Descriptions for Zero-shot Named Entity Recognition and Classification  
<https://aclanthology.org/2021.acl-long.120>  
**AUTHORS:** Rami Aly, Andreas Vlachos, Ryan McDonald  
**HIGHLIGHT:** This paper presents the first approach for zero-shot NERC, introducing novel architectures that leverage the fact that textual descriptions for many entity classes occur naturally.

121, **TITLE:** MECT: Multi-Metadata Embedding based Cross-Transformer for Chinese Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.121>  
**AUTHORS:** Shuang Wu, Xiaoning Song, Zhenhua Feng  
**HIGHLIGHT:** This paper presents a novel Multi-metadata Embedding based Cross-Transformer (MECT) to improve the performance of Chinese NER by fusing the structural information of Chinese characters.

122, **TITLE:** Factuality Assessment as Modal Dependency Parsing  
<https://aclanthology.org/2021.acl-long.122>  
**AUTHORS:** Jiarui Yao, Haoling Qiu, Jin Zhao, Bonan Min, Nianwen Xue  
**HIGHLIGHT:** In this paper, we put forward a novel idea of encoding the utterances with a modal dependency parsing task that identifies the events and their sources, formally known as conceivers, and then determine the level of certainty that the sources are asserting with respect to the events.

123, **TITLE:** Directed Acyclic Graph Network for Conversational Emotion Recognition  
<https://aclanthology.org/2021.acl-long.123>  
**AUTHORS:** Weizhou Shen, Siyue Wu, Yunyi Yang, Xiaojun Quan  
**HIGHLIGHT:** In this paper, we put forward a novel idea of encoding the utterances with a directed acyclic graph (DAG) to better model the intrinsic structure within a conversation, and design a directed acyclic neural network, namely DAG-ERC, to implement this idea.

124, **TITLE:** Improving Formality Style Transfer with Context-Aware Rule Injection  
<https://aclanthology.org/2021.acl-long.124>  
**AUTHORS:** Zonghai Yao, Hong Yu  
**HIGHLIGHT:** Here we present Context-Aware Rule Injection (CARI), an innovative method for formality style transfer (FST) by injecting multiple rules into an end-to-end BERT-based encoder and decoder model.

125, **TITLE:** Topic-Driven and Knowledge-Aware Transformer for Dialogue Emotion Detection  
<https://aclanthology.org/2021.acl-long.125>  
**AUTHORS:** Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, Yulan He  
**HIGHLIGHT:** In this paper, we propose a Topic-Driven Knowledge-Aware Transformer to handle the challenges above.

126, **TITLE:** Syntopical Graphs for Computational Argumentation Tasks  
<https://aclanthology.org/2021.acl-long.126>  
**AUTHORS:** Joe Barrow, Rajiv Jain, Nedim Lipka, Franck Dernoncourt, Vlad Morariu, Varun Manjunatha, Douglas Oard, Philip Resnik, Henning Wachsmuth  
**HIGHLIGHT:** We introduce a general approach to these tasks motivated by syntopical reading, a reading process that emphasizes comparing and contrasting viewpoints in order to improve topic understanding.

127, **TITLE:** Stance Detection in COVID-19 Tweets  
<https://aclanthology.org/2021.acl-long.127>  
**AUTHORS:** Kyle Glandt, Sarthak Khanal, Yingjie Li, Doina Caragea, Cornelia Caragea  
**HIGHLIGHT:** We annotate a new stance detection dataset, called COVID-19-Stance.

128, **TITLE:** Topic-Aware Evidence Reasoning and Stance-Aware Aggregation for Fact Verification

<https://aclanthology.org/2021.acl-long.128>

AUTHORS: Jiasheng Si, Deyu Zhou, Tongzhe Li, Xingyu Shi, Yulan He

HIGHLIGHT: To alleviate the above issues, we propose a novel topic-aware evidence reasoning and stance-aware aggregation model for more accurate fact verification, with the following four key properties: 1) checking topical consistency between the claim and evidence; 2) maintaining topical coherence among multiple pieces of evidence; 3) ensuring semantic similarity between the global topic information and the semantic representation of evidence; 4) aggregating evidence based on their implicit stances to the claim.

129, TITLE: Changes in European Solidarity Before and During COVID-19: Evidence from a Large Crowd- and Expert-Annotated Twitter Dataset

<https://aclanthology.org/2021.acl-long.129>

AUTHORS: Alexandra Ils, Dan Liu, Daniela Grunow, Steffen Eger

HIGHLIGHT: We introduce the well-established social scientific concept of social solidarity and its contestation, anti-solidarity, as a new problem setting to supervised machine learning in NLP to assess how European solidarity discourses changed before and after the COVID-19 outbreak was declared a global pandemic.

130, TITLE: Measuring Conversational Uptake: A Case Study on Student-Teacher Interactions

<https://aclanthology.org/2021.acl-long.130>

AUTHORS: Dorottya Demszky, Jing Liu, Zid Mancenido, Julie Cohen, Heather Hill, Dan Jurafsky, Tatsunori Hashimoto

HIGHLIGHT: We propose a framework for computationally measuring uptake, by (1) releasing a dataset of student-teacher exchanges extracted from US math classroom transcripts annotated for uptake by experts; (2) formalizing uptake as pointwise Jensen-Shannon Divergence (pJSD), estimated via next utterance classification; (3) conducting a linguistically-motivated comparison of different unsupervised measures and (4) correlating these measures with educational outcomes.

131, TITLE: A Survey of Code-switching: Linguistic and Social Perspectives for Language Technologies

<https://aclanthology.org/2021.acl-long.131>

AUTHORS: A. Seza Dogru?z, Sunayana Sitaram, Barbara E. Bullock, Almeida Jacqueline Toribio

HIGHLIGHT: To fill this gap, we offer a survey of code-switching (C-S) covering the literature in linguistics with a reflection on the key issues in language technologies.

132, TITLE: Learning from the Worst: Dynamically Generated Datasets to Improve Online Hate Detection

<https://aclanthology.org/2021.acl-long.132>

AUTHORS: Bertie Vidgen, Tristan Thrush, Zeerak Waseem, Douwe Kiela

HIGHLIGHT: We present a human-and-model-in-the-loop process for dynamically generating datasets and training better performing and more robust hate detection models.

133, TITLE: InfoSurgeon: Cross-Media Fine-grained Information Consistency Checking for Fake News Detection

<https://aclanthology.org/2021.acl-long.133>

AUTHORS: Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, Avi Sil

HIGHLIGHT: We contribute a novel benchmark for fake news detection at the knowledge element level, as well as a solution for this task which incorporates cross-media consistency checking to detect the fine-grained knowledge elements making news articles misinformative.

134, TITLE: I like fish, especially dolphins: Addressing Contradictions in Dialogue Modeling

<https://aclanthology.org/2021.acl-long.134>

AUTHORS: Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, Jason Weston

HIGHLIGHT: To quantify how well natural language understanding models can capture consistency in a general conversation, we introduce the DialogueE Contradiction DEtection task (DECODE) and a new conversational dataset containing both human-human and human-bot contradictory dialogues.

135, TITLE: A Sequence-to-Sequence Approach to Dialogue State Tracking

<https://aclanthology.org/2021.acl-long.135>

AUTHORS: Yue Feng, Yang Wang, Hang Li

HIGHLIGHT: This paper proposes a new approach to dialogue state tracking, referred to as Seq2Seq-DU, which formalizes DST as a sequence-to-sequence problem.

136, TITLE: Discovering Dialog Structure Graph for Coherent Dialog Generation

<https://aclanthology.org/2021.acl-long.136>

AUTHORS: Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che

HIGHLIGHT: To this end, we present an unsupervised model, Discrete Variational Auto-Encoder with Graph Neural Network (DVAE-GNN), to discover discrete hierarchical latent dialog states (at the level of both session and utterance) and their transitions from corpus as a dialog structure graph.

- 137, TITLE: Dialogue Response Selection with Hierarchical Curriculum Learning  
<https://aclanthology.org/2021.acl-long.137>  
AUTHORS: Yixuan Su, Deng Cai, Qingyu Zhou, Zibo Lin, Simon Baker, Yunbo Cao, Shuming Shi, Nigel Collier, Yan Wang  
HIGHLIGHT: Motivated by the recent finding that models trained with random negative samples are not ideal in real-world scenarios, we propose a hierarchical curriculum learning framework that trains the matching model in an "easy-to-difficult" scheme.
- 138, TITLE: A Joint Model for Dropped Pronoun Recovery and Conversational Discourse Parsing in Chinese Conversational Speech  
<https://aclanthology.org/2021.acl-long.138>  
AUTHORS: Jingxuan Yang, Kerui Xu, Jun Xu, Si Li, Sheng Gao, Jun Guo, Nianwen Xue, Ji-Rong Wen  
HIGHLIGHT: In this paper, we present a neural model for joint dropped pronoun recovery (DPR) and conversational discourse parsing (CDP) in Chinese conversational speech.
- 139, TITLE: A Systematic Investigation of KB-Text Embedding Alignment at Scale  
<https://aclanthology.org/2021.acl-long.139>  
AUTHORS: Vardaan Pahuja, Yu Gu, Wenhui Chen, Mehdi Bahrami, Lei Liu, Wei-Peng Chen, Yu Su  
HIGHLIGHT: We conduct a large-scale, systematic investigation of aligning KB and text embeddings for joint reasoning.
- 140, TITLE: Named Entity Recognition with Small Strongly Labeled and Large Weakly Labeled Data  
<https://aclanthology.org/2021.acl-long.140>  
AUTHORS: Haoming Jiang, Danqing Zhang, Tianyu Cao, Bing Yin, Tuo Zhao  
HIGHLIGHT: To address this issue, we propose a new multi-stage computational framework - NEEDLE with three essential ingredients: (1) weak label completion, (2) noise-aware loss function, and (3) final fine-tuning over the strongly labeled data.
- 141, TITLE: Ultra-Fine Entity Typing with Weak Supervision from a Masked Language Model  
<https://aclanthology.org/2021.acl-long.141>  
AUTHORS: Hongliang Dai, Yangqiu Song, Haixun Wang  
HIGHLIGHT: To remedy this problem, in this paper, we propose to obtain training data for ultra-fine entity typing by using a BERT Masked Language Model (MLM).
- 142, TITLE: Improving Named Entity Recognition by External Context Retrieving and Cooperative Learning  
<https://aclanthology.org/2021.acl-long.142>  
AUTHORS: Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, Kewei Tu  
HIGHLIGHT: In this paper, we propose to find external contexts of a sentence by retrieving and selecting a set of semantically relevant texts through a search engine, with the original sentence as the query.
- 143, TITLE: Implicit Representations of Meaning in Neural Language Models  
<https://aclanthology.org/2021.acl-long.143>  
AUTHORS: Belinda Z. Li, Maxwell Nye, Jacob Andreas  
HIGHLIGHT: In BART and T5 transformer language models, we identify contextual word representations that function as \*models of entities and situations\* as they evolve throughout a discourse.
- 144, TITLE: Causal Analysis of Syntactic Agreement Mechanisms in Neural Language Models  
<https://aclanthology.org/2021.acl-long.144>  
AUTHORS: Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, Yonatan Belinkov  
HIGHLIGHT: To elucidate the mechanisms by which the models accomplish this behavior, this study applies causal mediation analysis to pre-trained neural language models.
- 145, TITLE: Bird's Eye: Probing for Linguistic Graph Structures with a Simple Information-Theoretic Approach  
<https://aclanthology.org/2021.acl-long.145>  
AUTHORS: Yifan Hou, Mrinmaya Sachan  
HIGHLIGHT: In this paper, we propose a new information-theoretic probe, Bird's Eye, which is a fairly simple probe method for detecting if and how these representations encode the information in these linguistic graphs.
- 146, TITLE: Knowledgeable or Educated Guess? Revisiting Language Models as Knowledge Bases  
<https://aclanthology.org/2021.acl-long.146>  
AUTHORS: Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, Jin Xu

**HIGHLIGHT:** In this paper, we conduct a rigorous study to explore the underlying predicting mechanisms of MLMs over different extraction paradigms.

147, **TITLE:** Poisoning Knowledge Graph Embeddings via Relation Inference Patterns  
<https://aclanthology.org/2021.acl-long.147>

**AUTHORS:** Peru Bhardwaj, John Kelleher, Luca Costabello, Declan O'Sullivan

**HIGHLIGHT:** To poison KGE models, we propose to exploit their inductive abilities which are captured through the relationship patterns like symmetry, inversion and composition in the knowledge graph.

148, **TITLE:** Bad Seeds: Evaluating Lexical Methods for Bias Measurement  
<https://aclanthology.org/2021.acl-long.148>

**AUTHORS:** Maria Antoniak, David Mimno

**HIGHLIGHT:** We gather seeds used in prior work, documenting their common sources and rationales, and in case studies of three English-language corpora, we enumerate the different types of social biases and linguistic features that, once encoded in the seeds, can affect subsequent bias measurements.

149, **TITLE:** A Survey of Race, Racism, and Anti-Racism in NLP  
<https://aclanthology.org/2021.acl-long.149>

**AUTHORS:** Anjalie Field, Su Lin Blodgett, Zeerak Waseem, Yulia Tsvetkov

**HIGHLIGHT:** In this work, we survey 79 papers from the ACL anthology that mention race.

150, **TITLE:** Intrinsic Bias Metrics Do Not Correlate with Application Bias  
<https://aclanthology.org/2021.acl-long.150>

**AUTHORS:** Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sánchez, Mugdha Pandya, Adam Lopez

**HIGHLIGHT:** We compare intrinsic and extrinsic metrics across hundreds of trained models covering different tasks and experimental conditions.

151, **TITLE:** RedditBias: A Real-World Resource for Bias Evaluation and Debiasing of Conversational Language Models  
<https://aclanthology.org/2021.acl-long.151>

**AUTHORS:** Soumya Barikeri, Anne Lauscher, Ivan Vulic, Goran Glavač

**HIGHLIGHT:** In this work, we present REDDITBIAS, the first conversational data set grounded in the actual human conversations from Reddit, allowing for bias measurement and mitigation across four important bias dimensions: gender, race, religion, and queerness.

152, **TITLE:** Contributions of Transformer Attention Heads in Multi- and Cross-lingual Tasks  
<https://aclanthology.org/2021.acl-long.152>

**AUTHORS:** Weicheng Ma, Kai Zhang, Renze Lou, Lili Wang, Soroush Vosoughi

**HIGHLIGHT:** This paper studies the relative importance of attention heads in Transformer-based models to aid their interpretability in cross-lingual and multi-lingual tasks.

153, **TITLE:** Crafting Adversarial Examples for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.153>

**AUTHORS:** Xinze Zhang, Junzhe Zhang, Zhenhua Chen, Kun He

**HIGHLIGHT:** In this work, we investigate veritable evaluations of NMT adversarial attacks, and propose a novel method to craft NMT adversarial examples.

154, **TITLE:** UXLA: A Robust Unsupervised Data Augmentation Framework for Zero-Resource Cross-Lingual NLP  
<https://aclanthology.org/2021.acl-long.154>

**AUTHORS:** M Saiful Bari, Tasnim Mohiuddin, Shafiq Joty

**HIGHLIGHT:** We propose UXLA, a novel unsupervised data augmentation framework for zero-resource transfer learning scenarios.

155, **TITLE:** Glancing Transformer for Non-Autoregressive Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.155>

**AUTHORS:** Lihua Qian, Hao Zhou, Yu Bao, Mingxuan Wang, Lin Qiu, Weinan Zhang, Yong Yu, Lei Li

**HIGHLIGHT:** We propose the Glancing Language Model (GLM) for single-pass parallel generation models.

156, **TITLE:** Hierarchical Context-aware Network for Dense Video Event Captioning  
<https://aclanthology.org/2021.acl-long.156>

**AUTHORS:** Lei Ji, Xianglin Guo, Haoyang Huang, Xilin Chen



**HIGHLIGHT:** In this paper, we introduce a novel Hierarchical Context-aware Network for dense video event captioning (HCN) to capture context from various aspects.

157, **TITLE:** Control Image Captioning Spatially and Temporally

<https://aclanthology.org/2021.acl-long.157>

**AUTHORS:** Kun Yan, Lei Ji, Huaishao Luo, Ming Zhou, Nan Duan, Shuai Ma

**HIGHLIGHT:** This paper aims to solve this problem by proposing a novel model called LoopCAG, which connects Contrastive constraints and Attention Guidance in a Loop manner, engaged explicit spatial and temporal constraints to the generating process.

158, **TITLE:** Edited Media Understanding Frames: Reasoning About the Intent and Implications of Visual Misinformation

<https://aclanthology.org/2021.acl-long.158>

**AUTHORS:** Jeff Da, Maxwell Forbes, Rowan Zellers, Anthony Zheng, Jena D. Hwang, Antoine Bosselut, Yejin Choi

**HIGHLIGHT:** In this paper, we study Edited Media Frames, a new formalism to understand visual media manipulation as structured annotations with respect to the intents, emotional reactions, attacks on individuals, and the overall implications of disinformation.

159, **TITLE:** PIGLeT: Language Grounding Through Neuro-Symbolic Interaction in a 3D World

<https://aclanthology.org/2021.acl-long.159>

**AUTHORS:** Rowan Zellers, Ari Holtzman, Matthew Peters, Roozbeh Mottaghi, Aniruddha Kembhavi, Ali Farhadi, Yejin Choi

**HIGHLIGHT:** We propose PIGLeT: a model that learns physical commonsense knowledge through interaction, and then uses this knowledge to ground language.

160, **TITLE:** Modeling Fine-Grained Entity Types with Box Embeddings

<https://aclanthology.org/2021.acl-long.160>

**AUTHORS:** Yasumasa Onoe, Michael Boratko, Andrew McCallum, Greg Durrett

**HIGHLIGHT:** We study the ability of box embeddings, which embed concepts as d-dimensional hyperrectangles, to capture hierarchies of types even when these relationships are not defined explicitly in the ontology.

161, **TITLE:** ChineseBERT: Chinese Pretraining Enhanced by Glyph and Pinyin Information

<https://aclanthology.org/2021.acl-long.161>

**AUTHORS:** Zijun Sun, Xiaoya Li, Xiaofei Sun, Yuxian Meng, Xiang Ao, Qing He, Fei Wu, Jiwei Li

**HIGHLIGHT:** In this work, we propose ChineseBERT, which incorporates both the glyph and pinyin information of Chinese characters into language model pretraining.

162, **TITLE:** Weight Distillation: Transferring the Knowledge in Neural Network Parameters

<https://aclanthology.org/2021.acl-long.162>

**AUTHORS:** Ye Lin, Yanyang Li, Ziyang Wang, Bei Li, Quan Du, Tong Xiao, Jingbo Zhu

**HIGHLIGHT:** In this paper, we propose Weight Distillation to transfer the knowledge in parameters of a large neural network to a small neural network through a parameter generator.

163, **TITLE:** Optimizing Deeper Transformers on Small Datasets

<https://aclanthology.org/2021.acl-long.163>

**AUTHORS:** Peng Xu, Dhruv Kumar, Wei Yang, Wenjie Zi, Keyi Tang, Chenyang Huang, Jackie Chi Kit Cheung, Simon J.D. Prince, Yanshuai Cao

**HIGHLIGHT:** This work shows that this does not always need to be the case: with proper initialization and optimization, the benefits of very deep transformers can carry over to challenging tasks with small datasets, including Text-to-SQL semantic parsing and logical reading comprehension.

164, **TITLE:** BERTAC: Enhancing Transformer-based Language Models with Adversarially Pretrained Convolutional Neural Networks

<https://aclanthology.org/2021.acl-long.164>

**AUTHORS:** Jong-Hoon Oh, Ryu Iida, Julien Kloetzer, Kentaro Torisawa

**HIGHLIGHT:** Of course, we do not have answers now, but, as an attempt to find better neural architectures and training schemes, we pretrain a simple CNN using a GAN-style learning scheme and Wikipedia data, and then integrate it with standard TLMs.

165, **TITLE:** COVID-Fact: Fact Extraction and Verification of Real-World Claims on COVID-19 Pandemic

<https://aclanthology.org/2021.acl-long.165>

**AUTHORS:** Arkadiy Saakyan, Tuhin Chakrabarty, Smaranda Muresan

HIGHLIGHT: We introduce a FEVER-like dataset COVID-Fact of 4,086 claims concerning the COVID-19 pandemic.

166, TITLE: Explaining Relationships Between Scientific Documents  
<https://aclanthology.org/2021.acl-long.166>  
AUTHORS: Kelvin Luu, Xinyi Wu, Rik Koncel-Kedziorski, Kyle Lo, Isabel Cachola, Noah A. Smith  
HIGHLIGHT: We address the task of explaining relationships between two scientific documents using natural language text.

167, TITLE: IrEne: Interpretable Energy Prediction for Transformers  
<https://aclanthology.org/2021.acl-long.167>  
AUTHORS: Qingqing Cao, Yash Kumar Lal, Harsh Trivedi, Aruna Balasubramanian, Niranjan Balasubramanian  
HIGHLIGHT: We present IrEne, an interpretable and extensible energy prediction system that accurately predicts the inference energy consumption of a wide range of Transformer-based NLP models.

168, TITLE: Mitigating Bias in Session-based Cyberbullying Detection: A Non-Compromising Approach  
<https://aclanthology.org/2021.acl-long.168>  
AUTHORS: Lu Cheng, Ahmadreza Mosallanezhad, Yasin Silva, Deborah Hall, Huan Liu  
HIGHLIGHT: In this paper, we first show evidence of such bias in models trained on sessions collected from different social media platforms (e.g., Instagram). We then propose a context-aware and model-agnostic debiasing strategy that leverages a reinforcement learning technique, without requiring any extra resources or annotations apart from a pre-defined set of sensitive triggers commonly used for identifying cyberbullying instances.

169, TITLE: PlotCoder: Hierarchical Decoding for Synthesizing Visualization Code in Programmatic Context  
<https://aclanthology.org/2021.acl-long.169>  
AUTHORS: Xinyun Chen, Linyuan Gong, Alvin Cheung, Dawn Song  
HIGHLIGHT: In this paper, we propose the new task of synthesizing visualization programs from a combination of natural language utterances and code context.

170, TITLE: Changing the World by Changing the Data  
<https://aclanthology.org/2021.acl-long.170>  
AUTHORS: Anna Rogers  
HIGHLIGHT: This position paper maps out the arguments for and against data curation, and argues that fundamentally the point is moot: curation already is and will be happening, and it is changing the world.

171, TITLE: EarlyBERT: Efficient BERT Training via Early-bird Lottery Tickets  
<https://aclanthology.org/2021.acl-long.171>  
AUTHORS: Xiaohan Chen, Yu Cheng, Shuohang Wang, Zhe Gan, Zhangyang Wang, Jingjing Liu  
HIGHLIGHT: In this paper, inspired by the Early-Bird Lottery Tickets recently studied for computer vision tasks, we propose EarlyBERT, a general computationally-efficient training algorithm applicable to both pre-training and fine-tuning of large-scale language models.

172, TITLE: On the Effectiveness of Adapter-based Tuning for Pretrained Language Model Adaptation  
<https://aclanthology.org/2021.acl-long.172>  
AUTHORS: Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, Luo Si  
HIGHLIGHT: We first show that adapter-based tuning better mitigates forgetting issues than fine-tuning since it yields representations with less deviation from those generated by the initial PrLM.

173, TITLE: Data Augmentation for Text Generation Without Any Augmented Data  
<https://aclanthology.org/2021.acl-long.173>  
AUTHORS: Wei Bi, Huayang Li, Jiacheng Huang  
HIGHLIGHT: In this work, we derive an objective to formulate the problem of data augmentation on text generation tasks without any use of augmented data constructed by specific mapping functions.

174, TITLE: Integrating Semantics and Neighborhood Information with Graph-Driven Generative Models for Document Retrieval  
<https://aclanthology.org/2021.acl-long.174>  
AUTHORS: Zijiang Ou, Qinliang Su, Jianxing Yu, Bang Liu, Jingwen Wang, Ruihui Zhao, Changyou Chen, Yefeng Zheng  
HIGHLIGHT: In this paper, we encode the neighborhood information with a graph-induced Gaussian distribution, and propose to integrate the two types of information with a graph-driven generative model.

175, TITLE: SMURF: SeMantic and linguistic UndeRstanding Fusion for Caption Evaluation via Typicality Analysis

<https://aclanthology.org/2021.acl-long.175>

AUTHORS: Joshua Feinglass, Yezhou Yang

HIGHLIGHT: We introduce "typicality", a new formulation of evaluation rooted in information theory, which is uniquely suited for problems lacking a definite ground truth.

176, TITLE: KaggleDBQA: Realistic Evaluation of Text-to-SQL Parsers

<https://aclanthology.org/2021.acl-long.176>

AUTHORS: Chia-Hsuan Lee, Oleksandr Polozov, Matthew Richardson

HIGHLIGHT: In this work, we examine the challenges that still prevent these techniques from practical deployment.

177, TITLE: QASR: QCRI Aljazeera Speech Resource A Large Scale Annotated Arabic Speech Corpus

<https://aclanthology.org/2021.acl-long.177>

AUTHORS: Hamdy Mubarak, Amir Hussein, Shammur Absar Chowdhury, Ahmed Ali

HIGHLIGHT: We introduce the largest transcribed Arabic speech corpus, QASR, collected from the broadcast domain.

178, TITLE: An Empirical Study on Hyperparameter Optimization for Fine-Tuning Pre-trained Language Models

<https://aclanthology.org/2021.acl-long.178>

AUTHORS: Xueqing Liu, Chi Wang

HIGHLIGHT: In this paper, we investigate the performance of modern hyperparameter optimization methods (HPO) on fine-tuning pre-trained language models.

179, TITLE: Better than Average: Paired Evaluation of NLP systems

<https://aclanthology.org/2021.acl-long.179>

AUTHORS: Maxime Peyrard, Wei Zhao, Steffen Eger, Robert West

HIGHLIGHT: In this work, we question the use of averages for aggregating evaluation scores into a final number used to decide which system is best, since the average, as well as alternatives such as the median, ignores the pairing arising from the fact that systems are evaluated on the same test instances.

180, TITLE: Chase: A Large-Scale and Pragmatic Chinese Dataset for Cross-Database Context-Dependent Text-to-SQL

<https://aclanthology.org/2021.acl-long.180>

AUTHORS: Jiaqi Guo, Ziliang Si, Yu Wang, Qian Liu, Ming Fan, Jian-Guang Lou, Zijiang Yang, Ting Liu

HIGHLIGHT: In this work, we present Chase, a large-scale and pragmatic Chinese dataset for XDTS.

181, TITLE: CLINE: Contrastive Learning with Semantic Negative Examples for Natural Language Understanding

<https://aclanthology.org/2021.acl-long.181>

AUTHORS: Dong Wang, Ning Ding, Piji Li, Haitao Zheng

HIGHLIGHT: To address this problem, we propose Contrastive Learning with semantic Negative Examples (CLINE), which constructs semantic negative examples unsupervised to improve the robustness under semantically adversarial attacking.

182, TITLE: Tree-Structured Topic Modeling with Nonparametric Neural Variational Inference

<https://aclanthology.org/2021.acl-long.182>

AUTHORS: Ziyi Chen, Cheng Ding, Zusheng Zhang, Yanghui Rao, Haoran Xie

HIGHLIGHT: In this study, we develop a tree-structured topic model by leveraging nonparametric neural variational inference.

183, TITLE: ExCAR: Event Graph Knowledge Enhanced Explainable Causal Reasoning

<https://aclanthology.org/2021.acl-long.183>

AUTHORS: Li Du, Xiao Ding, Kai Xiong, Ting Liu, Bing Qin

HIGHLIGHT: To facilitate this, we present an Event graph knowledge enhanced explainable CAusal Reasoning framework (ExCAR).

184, TITLE: Distributed Representations of Emotion Categories in Emotion Space

<https://aclanthology.org/2021.acl-long.184>

AUTHORS: Xiangyu Wang, Chengqing Zong

HIGHLIGHT: In this article, we first propose a general framework to learn the distributed representations for emotion categories in emotion space from a given emotion classification dataset.

185, TITLE: Style is NOT a single variable: Case Studies for Cross-Stylistic Language Understanding

<https://aclanthology.org/2021.acl-long.185>

AUTHORS: Dongyeop Kang, Eduard Hovy

**HIGHLIGHT:** This paper provides the benchmark corpus (XSLUE) that combines existing datasets and collects a new one for sentence-level cross-style language understanding and evaluation.

186, **TITLE:** DynaSent: A Dynamic Benchmark for Sentiment Analysis

<https://aclanthology.org/2021.acl-long.186>

**AUTHORS:** Christopher Potts, Zhengxuan Wu, Atticus Geiger, Douwe Kiela

**HIGHLIGHT:** We introduce DynaSent ('Dynamic Sentiment'), a new English-language benchmark task for ternary (positive/negative/neutral) sentiment analysis.

187, **TITLE:** A Hierarchical VAE for Calibrating Attributes while Generating Text using Normalizing Flow

<https://aclanthology.org/2021.acl-long.187>

**AUTHORS:** Bidisha Samanta, Mohit Agrawal, Niloy Ganguly

**HIGHLIGHT:** In this paper we propose a hierarchical architecture for finer control over the attribute, preserving content using attribute dis-entanglement.

188, **TITLE:** A Unified Generative Framework for Aspect-based Sentiment Analysis

<https://aclanthology.org/2021.acl-long.188>

**AUTHORS:** Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, Zheng Zhang

**HIGHLIGHT:** In this paper, we redefine every subtask target as a sequence mixed by pointer indexes and sentiment class indexes, which converts all ABSA subtasks into a unified generative formulation.

189, **TITLE:** Discovering Dialogue Slots with Weak Supervision

<https://aclanthology.org/2021.acl-long.189>

**AUTHORS:** Vojtech Hudecek, Ondrej Dušek, Zhou Yu

**HIGHLIGHT:** We propose a method that eliminates this requirement: We use weak supervision from existing linguistic annotation models to identify potential slot candidates, then automatically identify domain-relevant slots by using clustering algorithms.

190, **TITLE:** Enhancing the generalization for Intent Classification and Out-of-Domain Detection in SLU

<https://aclanthology.org/2021.acl-long.190>

**AUTHORS:** Yilin Shen, Yen-Chang Hsu, Avik Ray, Hongxia Jin

**HIGHLIGHT:** This paper proposes to train a model with only IND data while supporting both IND intent classification and OOD detection.

191, **TITLE:** ProtAugment: Intent Detection Meta-Learning through Unsupervised Diverse Paraphrasing

<https://aclanthology.org/2021.acl-long.191>

**AUTHORS:** Thomas Dopierre, Christophe Gravier, Wilfried Logerais

**HIGHLIGHT:** In this work, we propose ProtAugment, a meta-learning algorithm for short texts classification (the intent detection task).

192, **TITLE:** Robustness Testing of Language Understanding in Task-Oriented Dialog

<https://aclanthology.org/2021.acl-long.192>

**AUTHORS:** Jiexi Liu, Ryuichi Takanobu, Jiaxin Wen, Dazhen Wan, Hongguang Li, Weiran Nie, Cheng Li, Wei Peng, Minlie Huang

**HIGHLIGHT:** In this paper, we conduct comprehensive evaluation and analysis with respect to the robustness of natural language understanding models, and introduce three important aspects related to language understanding in real-world dialog systems, namely, language variety, speech characteristics, and noise perturbation.

193, **TITLE:** Comprehensive Study: How the Context Information of Different Granularity Affects Dialogue State Tracking?

<https://aclanthology.org/2021.acl-long.193>

**AUTHORS:** Puhai Yang, Heyan Huang, Xian-Ling Mao

**HIGHLIGHT:** Thus, in this paper, we will study and discuss how the context information of different granularity affects dialogue state tracking.

194, **TITLE:** OTTers: One-turn Topic Transitions for Open-Domain Dialogue

<https://aclanthology.org/2021.acl-long.194>

**AUTHORS:** Karin Sevegnani, David M. Howcroft, Ioannis Konstas, Verena Rieser

**HIGHLIGHT:** We first collect a new dataset of human one-turn topic transitions, which we call OTTers. We then explore different strategies used by humans when asked to complete such a task, and notice that the use of a bridging utterance to connect the two topics is the approach used the most.

- 195, TITLE: Towards Robustness of Text-to-SQL Models against Synonym Substitution  
<https://aclanthology.org/2021.acl-long.195>  
AUTHORS: Yujian Gan, Xinyun Chen, Qiuping Huang, Matthew Purver, John R. Woodward, Jinxia Xie, Pengsheng Huang  
HIGHLIGHT: In this work, we investigate the robustness of text-to-SQL models to synonym substitution.
- 196, TITLE: KACE: Generating Knowledge Aware Contrastive Explanations for Natural Language Inference  
<https://aclanthology.org/2021.acl-long.196>  
AUTHORS: Qianglong Chen, Feng Ji, Xiangji Zeng, Feng-Lin Li, Ji Zhang, Haiqing Chen, Yin Zhang  
HIGHLIGHT: In this paper, we focus on generating contrastive explanations with counterfactual examples in NLI and propose a novel Knowledge-Aware Contrastive Explanation generation framework (KACE).
- 197, TITLE: Self-Guided Contrastive Learning for BERT Sentence Representations  
<https://aclanthology.org/2021.acl-long.197>  
AUTHORS: Taekuk Kim, Kang Min Yoo, Sang-goo Lee  
HIGHLIGHT: In this work, we propose a contrastive learning method that utilizes self-guidance for improving the quality of BERT sentence representations.
- 198, TITLE: LGESQL: Line Graph Enhanced Text-to-SQL Model with Mixed Local and Non-Local Relations  
<https://aclanthology.org/2021.acl-long.198>  
AUTHORS: Ruisheng Cao, Lu Chen, Zhi Chen, Yanbin Zhao, Su Zhu, Kai Yu  
HIGHLIGHT: This work aims to tackle the challenging heterogeneous graph encoding problem in the text-to-SQL task.
- 199, TITLE: Multi-stage Pre-training over Simplified Multimodal Pre-training Models  
<https://aclanthology.org/2021.acl-long.199>  
AUTHORS: Tongtong Liu, Fangxiang Feng, Xiaojie Wang  
HIGHLIGHT: In this paper, we propose a new Multi-stage Pre-training (MSP) method, which uses information at different granularities from word, phrase to sentence in both texts and images to pre-train a model in stages.
- 200, TITLE: Beyond Sentence-Level End-to-End Speech Translation: Context Helps  
<https://aclanthology.org/2021.acl-long.200>  
AUTHORS: Biao Zhang, Ivan Titov, Barry Haddow, Rico Sennrich  
HIGHLIGHT: We investigate several decoding approaches, and introduce in-model ensemble decoding which jointly performs document- and sentence-level translation using the same model.
- 201, TITLE: LayoutLMv2: Multi-modal Pre-training for Visually-rich Document Understanding  
<https://aclanthology.org/2021.acl-long.201>  
AUTHORS: Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, Min Zhang, Lidong Zhou  
HIGHLIGHT: We propose LayoutLMv2 architecture with new pre-training tasks to model the interaction among text, layout, and image in a single multi-modal framework.
- 202, TITLE: UNIMO: Towards Unified-Modal Understanding and Generation via Cross-Modal Contrastive Learning  
<https://aclanthology.org/2021.acl-long.202>  
AUTHORS: Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, Haifeng Wang  
HIGHLIGHT: In this work, we propose a UNified-MODal pre-training architecture, namely UNIMO, which can effectively adapt to both single-modal and multi-modal understanding and generation tasks.
- 203, TITLE: Missing Modality Imagination Network for Emotion Recognition with Uncertain Missing Modalities  
<https://aclanthology.org/2021.acl-long.203>  
AUTHORS: Jinming Zhao, Ruichen Li, Qin Jin  
HIGHLIGHT: In this work, we propose a unified model, Missing Modality Imagination Network (MMIN), to deal with the uncertain missing modality problem.
- 204, TITLE: Stacked Acoustic-and-Textual Encoding: Integrating the Pre-trained Models into Speech Translation Encoders  
<https://aclanthology.org/2021.acl-long.204>  
AUTHORS: Chen Xu, Bojie Hu, Yanyang Li, Yuhao Zhang, Shen Huang, Qi Ju, Tong Xiao, Jingbo Zhu  
HIGHLIGHT: In this work, we propose a Stacked Acoustic-and-Textual Encoding (SATE) method for speech translation.
- 205, TITLE: N-ary Constituent Tree Parsing with Recursive Semi-Markov Model

<https://aclanthology.org/2021.acl-long.205>

AUTHORS: Xin Xin, Jinlong Li, Zeqi Tan

HIGHLIGHT: In this paper, we study the task of graph-based constituent parsing in the setting that binarization is not conducted as a pre-processing step, where a constituent tree may consist of nodes with more than two children.

206, TITLE: Automated Concatenation of Embeddings for Structured Prediction

<https://aclanthology.org/2021.acl-long.206>

AUTHORS: Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, Kewei Tu

HIGHLIGHT: In this paper, we propose Automated Concatenation of Embeddings (ACE) to automate the process of finding better concatenations of embeddings for structured prediction tasks, based on a formulation inspired by recent progress on neural architecture search.

207, TITLE: Multi-View Cross-Lingual Structured Prediction with Minimum Supervision

<https://aclanthology.org/2021.acl-long.207>

AUTHORS: Zechuan Hu, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, Kewei Tu

HIGHLIGHT: In this paper, we propose a multi-view framework, by leveraging a small number of labeled target sentences, to effectively combine multiple source models into an aggregated source view at different granularity levels (language, sentence, or sub-structure), and transfer it to a target view based on a task-specific model.

208, TITLE: The Limitations of Limited Context for Constituency Parsing

<https://aclanthology.org/2021.acl-long.208>

AUTHORS: Yuchen Li, Andrej Risteski

HIGHLIGHT: In this work, we answer representational questions raised by the architectures in (Shen et al., 2018a, 2019), as well as some transition-based syntax-aware language models (Dyer et al., 2016): what kind of syntactic structure can current neural approaches to syntax represent?

209, TITLE: Neural Bi-Lexicalized PCFG Induction

<https://aclanthology.org/2021.acl-long.209>

AUTHORS: Songlin Yang, Yanpeng Zhao, Kewei Tu

HIGHLIGHT: In this paper, we propose an approach to parameterize L-PCFGs without making implausible independence assumptions.

210, TITLE: Riddit: Norms of Offensiveness for English Reddit Comments

<https://aclanthology.org/2021.acl-long.210>

AUTHORS: Rishav Hada, Sohi Sudhir, Pushkar Mishra, Helen Yannakoudakis, Saif M. Mohammad, Ekaterina Shutova

HIGHLIGHT: We create the first dataset of English language Reddit comments that has fine-grained, real-valued scores between -1 (maximally supportive) and 1 (maximally offensive).

211, TITLE: Towards Quantifiable Dialogue Coherence Evaluation

<https://aclanthology.org/2021.acl-long.211>

AUTHORS: Zheng Ye, Liucun Lu, Lishan Huang, Liang Lin, Xiaodan Liang

HIGHLIGHT: To address these limitations, we propose Quantifiable Dialogue Coherence Evaluation (QuantiDCE), a novel framework aiming to train a quantifiable dialogue coherence metric that can reflect the actual human rating standards.

212, TITLE: Assessing the Representations of Idiomaticity in Vector Models with a Noun Compound Dataset Labeled at Type and Token Levels

<https://aclanthology.org/2021.acl-long.212>

AUTHORS: Marcos Garcia, Tiago Kramer Vieira, Carolina Scarton, Marco Idiart, Aline Villavicencio

HIGHLIGHT: This paper presents the Noun Compound Type and Token Idiomaticity (NCTTI) dataset, with human annotations for 280 noun compounds in English and 180 in Portuguese at both type and token level.

213, TITLE: Factoring Statutory Reasoning as Language Understanding Challenges

<https://aclanthology.org/2021.acl-long.213>

AUTHORS: Nils Holzenberger, Benjamin Van Durme

HIGHLIGHT: Prior work introduced a resource that approached statutory reasoning as a monolithic textual entailment problem, with neural baselines performing nearly at-chance.

214, TITLE: Evaluating Evaluation Measures for Ordinal Classification and Ordinal Quantification

<https://aclanthology.org/2021.acl-long.214>

AUTHORS: Tetsuya Sakai

**HIGHLIGHT:** In the present study, we utilise data from the SemEval and NTCIR communities to clarify the properties of nine evaluation measures in the context of OC tasks, and six measures in the context of OQ tasks.

215, **TITLE:** Interpretable and Low-Resource Entity Matching via Decoupling Feature Learning from Decision Making  
<https://aclanthology.org/2021.acl-long.215>

**AUTHORS:** Zijun Yao, Chengjiang Li, Tiansi Dong, Xin Lv, Jifan Yu, Lei Hou, Juanzi Li, Yichi Zhang, Zelin Dai  
**HIGHLIGHT:** In this paper, we propose a novel EM framework that consists of Heterogeneous Information Fusion (HIF) and Key Attribute Tree (KAT) Induction to decouple feature representation from matching decision.

216, **TITLE:** Locate and Label: A Two-stage Identifier for Nested Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.216>

**AUTHORS:** Yongliang Shen, Xinyin Ma, Zeqi Tan, Shuai Zhang, Wen Wang, Weiming Lu  
**HIGHLIGHT:** First we generate span proposals by filtering and boundary regression on the seed spans to locate the entities, and then label the boundary-adjusted span proposals with the corresponding categories.

217, **TITLE:** Text2Event: Controllable Sequence-to-Structure Generation for End-to-end Event Extraction  
<https://aclanthology.org/2021.acl-long.217>

**AUTHORS:** Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, Shaoyi Chen  
**HIGHLIGHT:** In this paper, we propose Text2Event, a sequence-to-structure generation paradigm that can directly extract events from the text in an end-to-end manner.

218, **TITLE:** A Large-Scale Chinese Multimodal NER Dataset with Speech Clues  
<https://aclanthology.org/2021.acl-long.218>

**AUTHORS:** Dianbo Sui, Zhengkun Tian, Yubo Chen, Kang Liu, Jun Zhao  
**HIGHLIGHT:** In this paper, we aim to explore an uncharted territory, which is Chinese multimodal named entity recognition (NER) with both textual and acoustic contents.

219, **TITLE:** A Neural Transition-based Joint Model for Disease Named Entity Recognition and Normalization  
<https://aclanthology.org/2021.acl-long.219>

**AUTHORS:** Zongcheng Ji, Tian Xia, Mei Han, Jing Xiao  
**HIGHLIGHT:** In this work, we propose a neural transition-based joint model to alleviate these two issues.

220, **TITLE:** OntoED: Low-resource Event Detection with Ontology Embedding  
<https://aclanthology.org/2021.acl-long.220>

**AUTHORS:** Shumin Deng, Ningyu Zhang, Luoqi Li, Chen Hui, Tou Huaixiao, Mosha Chen, Fei Huang, Huajun Chen  
**HIGHLIGHT:** To address these problems, we formulate ED as a process of event ontology population: linking event instances to pre-defined event types in event ontology, and propose a novel ED framework entitled OntoED with ontology embedding.

221, **TITLE:** Self-Training Sampling with Monolingual Data Uncertainty for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.221>

**AUTHORS:** Wenxiang Jiao, Xing Wang, Zhaopeng Tu, Shuming Shi, Michael Lyu, Irwin King  
**HIGHLIGHT:** In this work, we propose to improve the sampling procedure by selecting the most informative monolingual sentences to complement the parallel data.

222, **TITLE:** Breaking the Corpus Bottleneck for Context-Aware Neural Machine Translation with Cross-Task Pre-training  
<https://aclanthology.org/2021.acl-long.222>

**AUTHORS:** Linqing Chen, Junhui Li, Zhengxian Gong, Boxing Chen, Weihua Luo, Min Zhang, Guodong Zhou  
**HIGHLIGHT:** To break the corpus bottleneck, in this paper we aim to improve context-aware NMT by taking the advantage of the availability of both large-scale sentence-level parallel dataset and source-side monolingual documents.

223, **TITLE:** Guiding Teacher Forcing with Seer Forcing for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.223>

**AUTHORS:** Yang Feng, Shuhao Gu, Dengji Guo, Zhengxin Yang, Chenze Shao  
**HIGHLIGHT:** To address this problem, we introduce another decoder, called seer decoder, into the encoder-decoder framework during training, which involves future information in target predictions.

224, **TITLE:** Cascade versus Direct Speech Translation: Do the Differences Still Make a Difference?  
<https://aclanthology.org/2021.acl-long.224>

**AUTHORS:** Luisa Bentivogli, Mauro Cettolo, Marco Gaido, Alina Karakanta, Alberto Martinelli, Matteo Negri, Marco Turchi

**HIGHLIGHT:** Starting from this question, we present a systematic comparison between state-of-the-art systems representative of the two paradigms.

225, **TITLE:** Unsupervised Neural Machine Translation for Low-Resource Domains via Meta-Learning  
<https://aclanthology.org/2021.acl-long.225>  
**AUTHORS:** Cheonbok Park, Yunwon Tae, TaeHee Kim, Soyoung Yang, Mohammad Azam Khan, Lucy Park, Jaegul Choo  
**HIGHLIGHT:** To address this issue, this paper presents a novel meta-learning algorithm for unsupervised neural machine translation (UNMT) that trains the model to adapt to another domain by utilizing only a small amount of training data.

226, **TITLE:** Lightweight Cross-Lingual Sentence Representation Learning  
<https://aclanthology.org/2021.acl-long.226>  
**AUTHORS:** Zhuoyuan Mao, Prakhar Gupta, Chenhui Chu, Martin Jaggi, Sadao Kurohashi  
**HIGHLIGHT:** In this work, we introduce a lightweight dual-transformer architecture with just 2 layers for generating memory-efficient cross-lingual sentence representations.

227, **TITLE:** ERNIE-Doc: A Retrospective Long-Document Modeling Transformer  
<https://aclanthology.org/2021.acl-long.227>  
**AUTHORS:** SiYu Ding, Junyuan Shang, Shuohuan Wang, Yu Sun, Hao Tian, Hua Wu, Haifeng Wang  
**HIGHLIGHT:** In this paper, we propose ERNIE-Doc, a document-level language pretraining model based on Recurrence Transformers.

228, **TITLE:** Marginal Utility Diminishes: Exploring the Minimum Knowledge for BERT Knowledge Distillation  
<https://aclanthology.org/2021.acl-long.228>  
**AUTHORS:** Yuanxin Liu, Fandong Meng, Zheng Lin, Weiping Wang, Jie Zhou  
**HIGHLIGHT:** In this way, we show that 1) the student's performance can be improved by extracting and distilling the crucial HSK, and 2) using a tiny fraction of HSK can achieve the same performance as extensive HSK distillation.

229, **TITLE:** Rational LAMOL: A Rationale-based Lifelong Learning Framework  
<https://aclanthology.org/2021.acl-long.229>  
**AUTHORS:** Kasidis Kanwatchara, Thanapapas Horsuwan, Piyawat Lertvittayakumjorn, Boonserm Kijisirikul, Peerapon Vateekul  
**HIGHLIGHT:** In this paper, we introduce Rational LAMOL, a novel end-to-end LL framework for language models.

230, **TITLE:** EnsLM: Ensemble Language Model for Data Diversity by Semantic Clustering  
<https://aclanthology.org/2021.acl-long.230>  
**AUTHORS:** Zhibin Duan, Hao Zhang, Chaojie Wang, Zhengjue Wang, Bo Chen, Mingyuan Zhou  
**HIGHLIGHT:** Having obtained the clustering assignment for each sample, we develop the ensemble LM (EnsLM) with the technique of weight modulation.

231, **TITLE:** LeeBERT: Learned Early Exit for BERT with cross-level optimization  
<https://aclanthology.org/2021.acl-long.231>  
**AUTHORS:** Wei Zhu  
**HIGHLIGHT:** In this work, to improve efficiency without performance drop, we propose a novel training scheme called Learned Early Exit for BERT (LeeBERT).

232, **TITLE:** Unsupervised Extractive Summarization-Based Representations for Accurate and Explainable Collaborative Filtering  
<https://aclanthology.org/2021.acl-long.232>  
**AUTHORS:** Reinald Adrian Pugoy, Hung-Yu Kao  
**HIGHLIGHT:** We pioneer the first extractive summarization-based collaborative filtering model called ESCOFILT.

233, **TITLE:** PLOME: Pre-training with Misspelled Knowledge for Chinese Spelling Correction  
<https://aclanthology.org/2021.acl-long.233>  
**AUTHORS:** Shulin Liu, Tao Yang, Tianchi Yue, Feng Zhang, Di Wang  
**HIGHLIGHT:** In this paper, we propose a Pre-trained masked Language model with Misspelled knowledge (PLOME) for CSC, which jointly learns how to understand language and correct spelling errors.

234, **TITLE:** Competence-based Multimodal Curriculum Learning for Medical Report Generation  
<https://aclanthology.org/2021.acl-long.234>  
**AUTHORS:** Fenglin Liu, Shen Ge, Xian Wu



**HIGHLIGHT:** To alleviate the data bias and make best use of available data, we propose a Competence-based Multimodal Curriculum Learning framework (CMCL).

235, **TITLE:** Learning Syntactic Dense Embedding with Correlation Graph for Automatic Readability Assessment  
<https://aclanthology.org/2021.acl-long.235>

**AUTHORS:** Xinying Qiu, Yuan Chen, Hanwu Chen, Jian-Yun Nie, Yuming Shen, Dawei Lu

**HIGHLIGHT:** We propose to incorporate linguistic features into neural network models by learning syntactic dense embeddings based on linguistic features.

236, **TITLE:** Meta-KD: A Meta Knowledge Distillation Framework for Language Model Compression across Domains  
<https://aclanthology.org/2021.acl-long.236>

**AUTHORS:** Haojie Pan, Chengyu Wang, Minghui Qiu, Yichang Zhang, Yaliang Li, Jun Huang

**HIGHLIGHT:** Hence we propose a Meta-Knowledge Distillation (Meta-KD) framework to build a meta-teacher model that captures transferable knowledge across domains and passes such knowledge to students.

237, **TITLE:** A Semantic-based Method for Unsupervised Commonsense Question Answering  
<https://aclanthology.org/2021.acl-long.237>

**AUTHORS:** Yilin Niu, Fei Huang, Jiaming Liang, Wenkai Chen, Xiaoyan Zhu, Minlie Huang

**HIGHLIGHT:** In this paper, we present a novel SEMantic-based Question Answering method (SEQA) for unsupervised commonsense question answering.

238, **TITLE:** Explanations for CommonsenseQA: New Dataset and Models

<https://aclanthology.org/2021.acl-long.238>

**AUTHORS:** Shourya Aggarwal, Divyanshu Mandowara, Vishwajeet Agrawal, Dinesh Khandelwal, Parag Singla, Dinesh Garg

**HIGHLIGHT:** Our explanation definition is based on certain desiderata, and translates an explanation into a set of positive and negative common-sense properties (aka facts) which not only explain the correct answer choice but also refute the incorrect ones.

239, **TITLE:** Few-Shot Question Answering by Pretraining Span Selection

<https://aclanthology.org/2021.acl-long.239>

**AUTHORS:** Ori Ram, Yuval Kirstain, Jonathan Berant, Amir Globerson, Omer Levy

**HIGHLIGHT:** We propose a new pretraining scheme tailored for question answering: recurring span selection.

240, **TITLE:** UnitedQA: A Hybrid Approach for Open Domain Question Answering

<https://aclanthology.org/2021.acl-long.240>

**AUTHORS:** Hao Cheng, Yelong Shen, Xiaodong Liu, Pengcheng He, Weizhu Chen, Jianfeng Gao

**HIGHLIGHT:** To date, most of recent work under the retrieval-reader framework for open-domain QA focuses on either extractive or generative reader exclusively. In this paper, we study a hybrid approach for leveraging the strengths of both models.

241, **TITLE:** Database reasoning over text

<https://aclanthology.org/2021.acl-long.241>

**AUTHORS:** James Thorne, Majid Yazdani, Marzieh Saeidi, Fabrizio Silvestri, Sebastian Riedel, Alon Halevy

**HIGHLIGHT:** We propose a modular architecture to answer these database-style queries over multiple spans from text and aggregating these at scale.

242, **TITLE:** Online Learning Meets Machine Translation Evaluation: Finding the Best Systems with the Least Human Effort

<https://aclanthology.org/2021.acl-long.242>

**AUTHORS:** V?nia Mendon?a, Ricardo Rei, Luisa Coheur, Alberto Sardinha, Ana L?cia Santos

**HIGHLIGHT:** To overcome the latter challenge, we propose a novel application of online learning that, given an ensemble of Machine Translation systems, dynamically converges to the best systems, by taking advantage of the human feedback available.

243, **TITLE:** How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models

<https://aclanthology.org/2021.acl-long.243>

**AUTHORS:** Phillip Rust, Jonas Pfeiffer, Ivan Vulic, Sebastian Ruder, Iryna Gurevych

**HIGHLIGHT:** In this work, we provide a systematic and comprehensive empirical comparison of pretrained multilingual language models versus their monolingual counterparts with regard to their monolingual task performance.

244, **TITLE:** Evaluating morphological typology in zero-shot cross-lingual transfer

<https://aclanthology.org/2021.acl-long.244>

**AUTHORS:** Antonio Mart?nez-Garc?a, Toni Badia, Jeremy Barnes

**HIGHLIGHT:** In this paper, we address what effects morphological typology has on zero-shot cross-lingual transfer for two tasks: Part-of-speech tagging and sentiment analysis.

245, **TITLE:** From Machine Translation to Code-Switching: Generating High-Quality Code-Switched Text  
<https://aclanthology.org/2021.acl-long.245>

**AUTHORS:** Ishan Tarunesh, Syamantak Kumar, Preethi Jyothi

**HIGHLIGHT:** In this work, we adapt a state-of-the-art neural machine translation model to generate Hindi-English code-switched sentences starting from monolingual Hindi sentences.

246, **TITLE:** Fast and Accurate Neural Machine Translation with Translation Memory  
<https://aclanthology.org/2021.acl-long.246>

**AUTHORS:** Qiuxiang He, Guoping Huang, Qu Cui, Li Li, Lema Liu

**HIGHLIGHT:** In this paper, we propose a fast and accurate approach to TM-based NMT within the Transformer framework: the model architecture is simple and employs a single bilingual sentence as its TM, leading to efficient training and inference; and its parameters are effectively optimized through a novel training criterion.

247, **TITLE:** Annotating Online Misogyny

<https://aclanthology.org/2021.acl-long.247>

**AUTHORS:** Philine Zeinert, Nanna Inie, Leon Derczynski

**HIGHLIGHT:** This paper makes three contributions in this area: Firstly, we describe the detailed design of our iterative annotation process and codebook. Secondly, we present a comprehensive taxonomy of labels for annotating misogyny in natural written language, and finally, we introduce a high-quality dataset of annotated posts sampled from social media posts.

248, **TITLE:** Few-NERD: A Few-shot Named Entity Recognition Dataset

<https://aclanthology.org/2021.acl-long.248>

**AUTHORS:** Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Haitao Zheng, Zhiyuan Liu

**HIGHLIGHT:** In this paper, we present Few-NERD, a large-scale human-annotated few-shot NER dataset with a hierarchy of 8 coarse-grained and 66 fine-grained entity types.

249, **TITLE:** MultiMET: A Multimodal Dataset for Metaphor Understanding

<https://aclanthology.org/2021.acl-long.249>

**AUTHORS:** Dongyu Zhang, Minghao Zhang, Heting Zhang, Liang Yang, Hongfei Lin

**HIGHLIGHT:** In this paper, we introduce MultiMET, a novel multimodal metaphor dataset to facilitate understanding metaphorical information from multimodal text and image.

250, **TITLE:** Human-in-the-Loop for Data Collection: a Multi-Target Counter Narrative Dataset to Fight Online Hate Speech  
<https://aclanthology.org/2021.acl-long.250>

**AUTHORS:** Margherita Fanton, Helena Bonaldi, Serra Sinem Tekiroglu, Marco Guerini

**HIGHLIGHT:** In this paper, we propose a novel human-in-the-loop data collection methodology in which a generative language model is refined iteratively by using its own data from the previous loops to generate new training samples that experts review and/or post-edit.

251, **TITLE:** Can Generative Pre-trained Language Models Serve As Knowledge Bases for Closed-book QA?

<https://aclanthology.org/2021.acl-long.251>

**AUTHORS:** Cunxiang Wang, Pai Liu, Yue Zhang

**HIGHLIGHT:** We construct a new dataset of closed-book QA using SQuAD, and investigate the performance of BART.

252, **TITLE:** Joint Models for Answer Verification in Question Answering Systems

<https://aclanthology.org/2021.acl-long.252>

**AUTHORS:** Zeyu Zhang, Thuy Vu, Alessandro Moschitti

**HIGHLIGHT:** This paper studies joint models for selecting correct answer sentences among the top k provided by answer sentence selection (AS2) modules, which are core components of retrieval-based Question Answering (QA) systems.

253, **TITLE:** Answering Ambiguous Questions through Generative Evidence Fusion and Round-Trip Prediction

<https://aclanthology.org/2021.acl-long.253>

**AUTHORS:** Yifan Gao, Henghui Zhu, Patrick Ng, Cicero Nogueira dos Santos, Zhiguo Wang, Feng Nan, Dejjiao Zhang, Ramesh Nallapati, Andrew O. Arnold, Bing Xiang

**HIGHLIGHT:** In this paper, we present a model that aggregates and combines evidence from multiple passages to adaptively predict a single answer or a set of question-answer pairs for ambiguous questions.

- 254, TITLE: TAT-QA: A Question Answering Benchmark on a Hybrid of Tabular and Textual Content in Finance  
<https://aclanthology.org/2021.acl-long.254>  
AUTHORS: Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, Tat-Seng Chua  
HIGHLIGHT: In this work, we extract samples from real financial reports to build a new large-scale QA dataset containing both Tabular And Textual data, named TAT-QA, where numerical reasoning is usually required to infer the answer, such as addition, subtraction, multiplication, division, counting, comparison/sorting, and the compositions.
- 255, TITLE: Modeling Transitions of Focal Entities for Conversational Knowledge Base Question Answering  
<https://aclanthology.org/2021.acl-long.255>  
AUTHORS: Yunshi Lan, Jing Jiang  
HIGHLIGHT: In this paper, we propose to model these implied entities, which we refer to as the focal entities of the conversation.
- 256, TITLE: Evidence-based Factual Error Correction  
<https://aclanthology.org/2021.acl-long.256>  
AUTHORS: James Thorne, Andreas Vlachos  
HIGHLIGHT: This paper introduces the task of factual error correction: performing edits to a claim so that the generated rewrite is better supported by evidence.
- 257, TITLE: Probabilistic, Structure-Aware Algorithms for Improved Variety, Accuracy, and Coverage of AMR Alignments  
<https://aclanthology.org/2021.acl-long.257>  
AUTHORS: Austin Blodgett, Nathan Schneider  
HIGHLIGHT: We present algorithms for aligning components of Abstract Meaning Representation (AMR) graphs to spans in English sentences.
- 258, TITLE: Meta-Learning to Compositionally Generalize  
<https://aclanthology.org/2021.acl-long.258>  
AUTHORS: Henry Conklin, Bailin Wang, Kenny Smith, Ivan Titov  
HIGHLIGHT: We implement a meta-learning augmented version of supervised learning whose objective directly optimizes for out-of-distribution generalization.
- 259, TITLE: Taming Pre-trained Language Models with N-gram Representations for Low-Resource Domain Adaptation  
<https://aclanthology.org/2021.acl-long.259>  
AUTHORS: Shizhe Diao, Ruijia Xu, Hongjin Su, Yilei Jiang, Yan Song, Tong Zhang  
HIGHLIGHT: In this paper, we aim to adapt a generic pretrained model with a relatively small amount of domain-specific data.
- 260, TITLE: ERICA: Improving Entity and Relation Understanding for Pre-trained Language Models via Contrastive Learning  
<https://aclanthology.org/2021.acl-long.260>  
AUTHORS: Yujia Qin, Yankai Lin, Ryuichi Takanobu, Zhiyuan Liu, Peng Li, Heng Ji, Minlie Huang, Maosong Sun, Jie Zhou  
HIGHLIGHT: To address this issue, we propose a novel contrastive learning framework ERICA to obtain a deep understanding of the entities and their relations in text.
- 261, TITLE: Position Bias Mitigation: A Knowledge-Aware Graph Model for Emotion Cause Extraction  
<https://aclanthology.org/2021.acl-long.261>  
AUTHORS: Hanqi Yan, Lin Gui, Gabriele Pergola, Yulan He  
HIGHLIGHT: To address the dataset bias, we propose a novel graph-based method to explicitly model the emotion triggering paths by leveraging the commonsense knowledge to enhance the semantic dependencies between a candidate clause and an emotion clause.
- 262, TITLE: Every Bite Is an Experience: Key Point Analysis of Business Reviews  
<https://aclanthology.org/2021.acl-long.262>  
AUTHORS: Roy Bar-Haim, Lilach Eden, Yoav Kantor, Roni Friedman, Noam Slonim  
HIGHLIGHT: We adapt KPA to review data by introducing Collective Key Point Mining for better key point extraction; integrating sentiment analysis into KPA; identifying good key point candidates for review summaries; and leveraging the massive amount of available reviews and their metadata.
- 263, TITLE: Structured Sentiment Analysis as Dependency Graph Parsing

- <https://aclanthology.org/2021.acl-long.263>  
AUTHORS: Jeremy Barnes, Robin Kurtz, Stephan Oepen, Lilja Vrelid, Erik Velldal  
HIGHLIGHT: We argue that this division has become counterproductive and propose a new unified framework to remedy the situation.
- 264, TITLE: Consistency Regularization for Cross-Lingual Fine-Tuning  
<https://aclanthology.org/2021.acl-long.264>  
AUTHORS: Bo Zheng, Li Dong, Shaohan Huang, Wenhui Wang, Zewen Chi, Saksham Singhal, Wanxiang Che, Ting Liu, Xia Song, Furu Wei  
HIGHLIGHT: In this work, we propose to improve cross-lingual fine-tuning with consistency regularization.
- 265, TITLE: Improving Pretrained Cross-Lingual Language Models via Self-Labeled Word Alignment  
<https://aclanthology.org/2021.acl-long.265>  
AUTHORS: Zewen Chi, Li Dong, Bo Zheng, Shaohan Huang, Xian-Ling Mao, Heyan Huang, Furu Wei  
HIGHLIGHT: In this paper, we introduce denoising word alignment as a new cross-lingual pre-training task.
- 266, TITLE: Rejuvenating Low-Frequency Words: Making the Most of Parallel Data in Non-Autoregressive Translation  
<https://aclanthology.org/2021.acl-long.266>  
AUTHORS: Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, Zhaopeng Tu  
HIGHLIGHT: Accordingly, we propose reverse KD to rejuvenate more alignments for low-frequency target words.
- 267, TITLE: G-Transformer for Document-Level Machine Translation  
<https://aclanthology.org/2021.acl-long.267>  
AUTHORS: Guangsheng Bao, Yue Zhang, Zhiyang Teng, Boxing Chen, Weihua Luo  
HIGHLIGHT: As a solution, we propose G-Transformer, introducing locality assumption as an inductive bias into Transformer, reducing the hypothesis space of the attention from target to source.
- 268, TITLE: Prevent the Language Model from being Overconfident in Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.268>  
AUTHORS: Mengqi Miao, Fandong Meng, Yijin Liu, Xiao-Hua Zhou, Jie Zhou  
HIGHLIGHT: Based on the property, we propose a Margin-based Token-level Objective (MTO) and a Margin-based Sentence-level Objective (MSO) to maximize the Margin for preventing the LM from being overconfident.
- 269, TITLE: Towards Emotional Support Dialog Systems  
<https://aclanthology.org/2021.acl-long.269>  
AUTHORS: Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, Minlie Huang  
HIGHLIGHT: In this paper, we define the Emotional Support Conversation (ESC) task and propose an ESC Framework, which is grounded on the Helping Skills Theory.
- 270, TITLE: Novel Slot Detection: A Benchmark for Discovering Unknown Slot Types in the Task-Oriented Dialogue System  
<https://aclanthology.org/2021.acl-long.270>  
AUTHORS: Yanan Wu, Zhiyuan Zeng, Keqing He, Hong Xu, Yuanmeng Yan, Huixing Jiang, Weiran Xu  
HIGHLIGHT: In this paper, we introduce a new task, Novel Slot Detection (NSD), in the task-oriented dialogue system.
- 271, TITLE: GTM: A Generative Triple-wise Model for Conversational Question Generation  
<https://aclanthology.org/2021.acl-long.271>  
AUTHORS: Lei Shen, Fandong Meng, Jinchao Zhang, Yang Feng, Jie Zhou  
HIGHLIGHT: To tackle these problems, we propose a generative triple-wise model with hierarchical variations for open-domain conversational question generation (CQG).
- 272, TITLE: Diversifying Dialog Generation via Adaptive Label Smoothing  
<https://aclanthology.org/2021.acl-long.272>  
AUTHORS: Yida Wang, Yinhe Zheng, Yong Jiang, Minlie Huang  
HIGHLIGHT: In this paper, we propose an Adaptive Label Smoothing (AdaLabel) approach that can adaptively estimate a target label distribution at each time step for different contexts.
- 273, TITLE: Out-of-Scope Intent Detection with Self-Supervision and Discriminative Training  
<https://aclanthology.org/2021.acl-long.273>  
AUTHORS: Li-Ming Zhan, Haowen Liang, Bo Liu, Lu Fan, Xiao-Ming Wu, Albert Y.S. Lam

**HIGHLIGHT:** In this paper, we propose a simple yet effective method to train an out-of-scope intent classifier in a fully end-to-end manner by simulating the test scenario in training, which requires no assumption on data distribution and no additional post-processing or threshold setting.

274, **TITLE:** Document-level Event Extraction via Heterogeneous Graph-based Interaction Model with a Tracker  
<https://aclanthology.org/2021.acl-long.274>

**AUTHORS:** Runxin Xu, Tianyu Liu, Lei Li, Baobao Chang

**HIGHLIGHT:** In this paper, we propose Heterogeneous Graph-based Interaction Model with a Tracker (GIT) to solve the aforementioned two challenges.

275, **TITLE:** Nested Named Entity Recognition via Explicitly Excluding the Influence of the Best Path  
<https://aclanthology.org/2021.acl-long.275>

**AUTHORS:** Yiran Wang, Hiroyuki Shindo, Yuji Matsumoto, Taro Watanabe

**HIGHLIGHT:** This paper presents a novel method for nested named entity recognition.

276, **TITLE:** LearnDA: Learnable Knowledge-Guided Data Augmentation for Event Causality Identification  
<https://aclanthology.org/2021.acl-long.276>

**AUTHORS:** Xinyu Zuo, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Weihua Peng, Yuguang Chen

**HIGHLIGHT:** To solve the data lacking problem, we introduce a new approach to augment training data for event causality identification, by iteratively generating new examples and classifying event causality in a dual learning framework.

277, **TITLE:** Revisiting the Negative Data of Distantly Supervised Relation Extraction

<https://aclanthology.org/2021.acl-long.277>

**AUTHORS:** Chenhao Xie, Jiaqing Liang, Jingping Liu, Chengsong Huang, Wenhao Huang, Yanghua Xiao

**HIGHLIGHT:** In this paper, we first provide a thorough analysis of the above challenges caused by negative data. Next, we formulate the problem of relation extraction into as a positive unlabeled learning task to alleviate false negative problem.

278, **TITLE:** Knowing the No-match: Entity Alignment with Dangling Cases

<https://aclanthology.org/2021.acl-long.278>

**AUTHORS:** Zequn Sun, Muhao Chen, Wei Hu

**HIGHLIGHT:** This paper studies a new problem setting of entity alignment for knowledge graphs (KGs).

279, **TITLE:** Superbizarre Is Not Superb: Derivational Morphology Improves BERT's Interpretation of Complex Words

<https://aclanthology.org/2021.acl-long.279>

**AUTHORS:** Valentin Hofmann, Janet Pierrehumbert, Hinrich Sch?tze

**HIGHLIGHT:** We present the first study investigating this question, taking BERT as the example PLM and focusing on its semantic representations of English derivatives.

280, **TITLE:** BERT is to NLP what AlexNet is to CV: Can Pre-Trained Language Models Identify Analogies?

<https://aclanthology.org/2021.acl-long.280>

**AUTHORS:** Asahi Ushio, Luis Espinosa Anke, Steven Schockaert, Jose Camacho-Collados

**HIGHLIGHT:** In this paper, we analyze the capabilities of transformer-based language models on this unsupervised task, using benchmarks obtained from educational settings, as well as more commonly used datasets.

281, **TITLE:** Exploring the Representation of Word Meanings in Context: A Case Study on Homonymy and Synonymy

<https://aclanthology.org/2021.acl-long.281>

**AUTHORS:** Marcos Garcia

**HIGHLIGHT:** This paper presents a multilingual study of word meaning representations in context.

282, **TITLE:** Measuring Fine-Grained Domain Relevance of Terms: A Hierarchical Core-Fringe Approach

<https://aclanthology.org/2021.acl-long.282>

**AUTHORS:** Jie Huang, Kevin Chang, JinJun Xiong, Wen-mei Hwu

**HIGHLIGHT:** We propose to measure fine-grained domain relevance- the degree that a term is relevant to a broad (e.g., computer science) or narrow (e.g., deep learning) domain.

283, **TITLE:** HERALD: An Annotation Efficient Method to Detect User Disengagement in Social Conversations

<https://aclanthology.org/2021.acl-long.283>

**AUTHORS:** Weixin Liang, Kai-Hui Liang, Zhou Yu

**HIGHLIGHT:** We propose HERALD, an efficient annotation framework that reframes the training data annotation process as a denoising problem.

- 284, TITLE: Value-Agnostic Conversational Semantic Parsing  
<https://aclanthology.org/2021.acl-long.284>  
AUTHORS: Emmanouil Antonios Platanios, Adam Pauls, Subhro Roy, Yuchen Zhang, Alexander Kyte, Alan Guo, Sam Thomson, Jayant Krishnamurthy, Jason Wolfe, Jacob Andreas, Dan Klein  
HIGHLIGHT: We propose a model that abstracts over values to focus prediction on type- and function-level context.
- 285, TITLE: MPC-BERT: A Pre-Trained Language Model for Multi-Party Conversation Understanding  
<https://aclanthology.org/2021.acl-long.285>  
AUTHORS: Jia-Chen Gu, Chongyang Tao, Zhenhua Ling, Can Xu, Xiubo Geng, Daxin Jiang  
HIGHLIGHT: To this end, we present MPC-BERT, a pre-trained model for MPC understanding that considers learning who says what to whom in a unified model with several elaborated self-supervised tasks.
- 286, TITLE: Best of Both Worlds: Making High Accuracy Non-incremental Transformer-based Disfluency Detection Incremental  
<https://aclanthology.org/2021.acl-long.286>  
AUTHORS: Morteza Rohanian, Julian Hough  
HIGHLIGHT: We address the challenge of introducing methods for word-by-word left-to-right incremental processing to Transformers such as BERT, models without an intrinsic sense of linear order.
- 287, TITLE: NeuralWOZ: Learning to Collect Task-Oriented Dialogue via Model-Based Simulation  
<https://aclanthology.org/2021.acl-long.287>  
AUTHORS: Sungdong Kim, Minsuk Chang, Sang-Woo Lee  
HIGHLIGHT: We propose NeuralWOZ, a novel dialogue collection framework that uses model-based dialogue simulation.
- 288, TITLE: CDRNN: Discovering Complex Dynamics in Human Language Processing  
<https://aclanthology.org/2021.acl-long.288>  
AUTHORS: Cory Shain  
HIGHLIGHT: This study proposes the continuous-time deconvolutional regressive neural network (CDRNN), a deep neural extension of continuous-time deconvolutional regression (Shain & Schuler, 2021) that jointly captures time-varying, non-linear, and delayed influences of predictors (e.g. word surprisal) on the response (e.g. reading time).
- 289, TITLE: Structural Guidance for Transformer Language Models  
<https://aclanthology.org/2021.acl-long.289>  
AUTHORS: Peng Qian, Tahira Naseem, Roger Levy, Ramon Fernandez Astudillo  
HIGHLIGHT: Here we study whether structural guidance leads to more human-like systematic linguistic generalization in Transformer language models without resorting to pre-training on very large amounts of data.
- 290, TITLE: Surprisal Estimators for Human Reading Times Need Character Models  
<https://aclanthology.org/2021.acl-long.290>  
AUTHORS: Byung-Doh Oh, Christian Clark, William Schuler  
HIGHLIGHT: This paper presents a character model that can be applied to a structural parser-based processing model to calculate word generation probabilities.
- 291, TITLE: CogAlign: Learning to Align Textual Neural Representations to Cognitive Language Processing Signals  
<https://aclanthology.org/2021.acl-long.291>  
AUTHORS: Yuqi Ren, Deyi Xiong  
HIGHLIGHT: In this paper, we propose a CogAlign approach to these issues, which learns to align textual neural representations to cognitive features.
- 292, TITLE: Self-Attention Networks Can Process Bounded Hierarchical Languages  
<https://aclanthology.org/2021.acl-long.292>  
AUTHORS: Shunyu Yao, Binghui Peng, Christos Papadimitriou, Karthik Narasimhan  
HIGHLIGHT: Specifically, we construct a hard-attention network with  $D+1$  layers and  $O(\log k)$  memory size (per token per layer) that recognizes Dyck- $(k, D)$ , and a soft-attention network with two layers and  $O(\log k)$  memory size that generates Dyck- $(k, D)$ .
- 293, TITLE: TextSETTR: Few-Shot Text Style Extraction and Tunable Targeted Restyling  
<https://aclanthology.org/2021.acl-long.293>  
AUTHORS: Parker Riley, Noah Constant, Mandy Guo, Girish Kumar, David Uthus, Zarana Parekh  
HIGHLIGHT: We present a novel approach to the problem of text style transfer.

- 294, TITLE: H-Transformer-1D: Fast One-Dimensional Hierarchical Attention for Sequences  
<https://aclanthology.org/2021.acl-long.294>  
AUTHORS: Zhenhai Zhu, Radu Soricut  
HIGHLIGHT: We describe an efficient hierarchical method to compute attention in the Transformer architecture.
- 295, TITLE: Making Pre-trained Language Models Better Few-shot Learners  
<https://aclanthology.org/2021.acl-long.295>  
AUTHORS: Tianyu Gao, Adam Fisch, Danqi Chen  
HIGHLIGHT: We present LM-BFF-better few-shot fine-tuning of language models-a suite of simple and complementary techniques for fine-tuning language models on a small number of annotated examples.
- 296, TITLE: A Sweet Rabbit Hole by DARCY: Using Honeypots to Detect Universal Trigger's Adversarial Attacks  
<https://aclanthology.org/2021.acl-long.296>  
AUTHORS: Thai Le, Noseong Park, Dongwon Lee  
HIGHLIGHT: To defend against this attack that can cause significant harm, in this paper, we borrow the "honeypot" concept from the cybersecurity community and propose DARCY, a honeypot-based defense framework against UniTrigger.
- 297, TITLE: Towards Propagation Uncertainty: Edge-enhanced Bayesian Graph Convolutional Networks for Rumor Detection  
<https://aclanthology.org/2021.acl-long.297>  
AUTHORS: Lingwei Wei, Dou Hu, Wei Zhou, Zhaojuan Yue, Songlin Hu  
HIGHLIGHT: Specifically, we propose a novel Edge-enhanced Bayesian Graph Convolutional Network (EBGCN) to capture robust structural features.
- 298, TITLE: Label-Specific Dual Graph Neural Network for Multi-Label Text Classification  
<https://aclanthology.org/2021.acl-long.298>  
AUTHORS: Qianwen Ma, Chunyuan Yuan, Wei Zhou, Songlin Hu  
HIGHLIGHT: In this paper, we propose a novel label-specific dual graph neural network (LDGN), which incorporates category information to learn label-specific components from documents, and employs dual Graph Convolution Network (GCN) to model complete and adaptive interactions among these components based on the statistical label co-occurrence and dynamic reconstruction graph in a joint way.
- 299, TITLE: TAN-NTM: Topic Attention Networks for Neural Topic Modeling  
<https://aclanthology.org/2021.acl-long.299>  
AUTHORS: Madhur Panwar, Shashank Shailabh, Milan Aggarwal, Balaji Krishnamurthy  
HIGHLIGHT: We propose a novel attention mechanism which factors in topic-word distribution to enable the model to attend on relevant words that convey topic related cues.
- 300, TITLE: Cross-language Sentence Selection via Data Augmentation and Rationale Training  
<https://aclanthology.org/2021.acl-long.300>  
AUTHORS: Yanda Chen, Chris Kedzie, Suraj Nair, Petra Galuscakova, Rui Zhang, Douglas Oard, Kathleen McKeown  
HIGHLIGHT: This paper proposes an approach to cross-language sentence selection in a low-resource setting.
- 301, TITLE: A Neural Model for Joint Document and Snippet Ranking in Question Answering for Large Document Collections  
<https://aclanthology.org/2021.acl-long.301>  
AUTHORS: Dimitris Pappas, Ion Androutsopoulos  
HIGHLIGHT: We present an architecture for joint document and snippet ranking, the two middle stages, which leverages the intuition that relevant documents have good snippets and good snippets come from relevant documents.
- 302, TITLE: W-RST: Towards a Weighted RST-style Discourse Framework  
<https://aclanthology.org/2021.acl-long.302>  
AUTHORS: Patrick Huber, Wen Xiao, Giuseppe Carenini  
HIGHLIGHT: Aiming for a better integration of data-driven and linguistically-inspired approaches, we explore whether RST Nuclearity, assigning a binary assessment of importance between text segments, can be replaced by automatically generated, real-valued scores, in what we call a Weighted-RST framework.
- 303, TITLE: ABCD: A Graph Framework to Convert Complex Sentences to a Covering Set of Simple Sentences  
<https://aclanthology.org/2021.acl-long.303>

- AUTHORS: Yanjun Gao, Ting-Hao Huang, Rebecca J. Passonneau  
HIGHLIGHT: We propose a new task to decompose each complex sentence into simple sentences derived from the tensed clauses in the source, and a novel problem formulation as a graph edit task.
- 304, TITLE: Which Linguist Invented the Lightbulb? Presupposition Verification for Question-Answering  
<https://aclanthology.org/2021.acl-long.304>  
AUTHORS: Najoung Kim, Ellie Pavlick, Burcu Karagol Ayan, Deepak Ramachandran  
HIGHLIGHT: Through a user preference study, we demonstrate that the oracle behavior of our proposed system-which provides responses based on presupposition failure-is preferred over the oracle behavior of existing QA systems. Then, we present a novel framework for implementing such a system in three steps: presupposition generation, presupposition verification, and explanation generation, reporting progress on each.
- 305, TITLE: Adversarial Learning for Discourse Rhetorical Structure Parsing  
<https://aclanthology.org/2021.acl-long.305>  
AUTHORS: Longyin Zhang, Fang Kong, Guodong Zhou  
HIGHLIGHT: In this work, we present our insight on evaluating the pros and cons of the entire DRS tree for global optimization.
- 306, TITLE: Exploring Discourse Structures for Argument Impact Classification  
<https://aclanthology.org/2021.acl-long.306>  
AUTHORS: Xin Liu, Jiefu Ou, Yangqiu Song, Xin Jiang  
HIGHLIGHT: This paper empirically shows that the discourse relations between two arguments along the context path are essential factors for identifying the persuasive power of an argument.
- 307, TITLE: Point, Disambiguate and Copy: Incorporating Bilingual Dictionaries for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.307>  
AUTHORS: Tong Zhang, Long Zhang, Wei Ye, Bo Li, Jinan Sun, Xiaoyu Zhu, Wen Zhao, Shikun Zhang  
HIGHLIGHT: This paper proposes a sophisticated neural architecture to incorporate bilingual dictionaries into Neural Machine Translation (NMT) models.
- 308, TITLE: VECO: Variable and Flexible Cross-lingual Pre-training for Language Understanding and Generation  
<https://aclanthology.org/2021.acl-long.308>  
AUTHORS: Fuli Luo, Wei Wang, Jiahao Liu, Yijia Liu, Bin Bi, Songfang Huang, Fei Huang, Luo Si  
HIGHLIGHT: In this paper, we plug a cross-attention module into the Transformer encoder to explicitly build the interdependence between languages.
- 309, TITLE: A unified approach to sentence segmentation of punctuated text in many languages  
<https://aclanthology.org/2021.acl-long.309>  
AUTHORS: Rachel Wicks, Matt Post  
HIGHLIGHT: We introduce a modern context-based modeling approach that provides a solution to the problem of segmenting punctuated text in many languages, and show how it can be trained on noisily-annotated data.
- 310, TITLE: Towards User-Driven Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.310>  
AUTHORS: Huan Lin, Liang Yao, Baosong Yang, Dayiheng Liu, Haibo Zhang, Weihua Luo, Degen Huang, Jinsong Su  
HIGHLIGHT: To fill this gap, we introduce a novel framework called user-driven NMT.
- 311, TITLE: End-to-End Lexically Constrained Machine Translation for Morphologically Rich Languages  
<https://aclanthology.org/2021.acl-long.311>  
AUTHORS: Josef Jon, Jo?o Paulo Aires, Dusan Varis, Ondrej Bojar  
HIGHLIGHT: In particular, we focus on methods based on training the model with constraints provided as part of the input sequence.
- 312, TITLE: Handling Extreme Class Imbalance in Technical Logbook Datasets  
<https://aclanthology.org/2021.acl-long.312>  
AUTHORS: Farhad Akhbardeh, Cecilia Ovesdotter Alm, Marcos Zampieri, Travis Desell  
HIGHLIGHT: In this paper we focus on the problem of technical issue classification by considering logbook datasets from the automotive, aviation, and facilities maintenance domains.
- 313, TITLE: ILDC for CJPE: Indian Legal Documents Corpus for Court Judgment Prediction and Explanation



<https://aclanthology.org/2021.acl-long.313>

AUTHORS: Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, Ashutosh Modi  
HIGHLIGHT: To promote research in developing such a system, we introduce ILDC (Indian Legal Documents Corpus).

314, TITLE: Supporting Cognitive and Emotional Empathic Writing of Students

<https://aclanthology.org/2021.acl-long.314>

AUTHORS: Thiemo Wambsganss, Christina Niklaus, Matthias S?llner, Siegfried Handschuh, Jan Marco Leimeister  
HIGHLIGHT: We present an annotation approach to capturing emotional and cognitive empathy in student-written peer reviews on business models in German.

315, TITLE: Dual Reader-Parser on Hybrid Textual and Tabular Evidence for Open Domain Question Answering

<https://aclanthology.org/2021.acl-long.315>

AUTHORS: Alexander Hanbo Li, Patrick Ng, Peng Xu, Henghui Zhu, Zhiguo Wang, Bing Xiang  
HIGHLIGHT: In this paper, we propose a hybrid framework that takes both textual and tabular evidences as input and generates either direct answers or SQL queries depending on which form could better answer the question.

316, TITLE: Generation-Augmented Retrieval for Open-Domain Question Answering

<https://aclanthology.org/2021.acl-long.316>

AUTHORS: Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, Weizhu Chen  
HIGHLIGHT: We propose Generation-Augmented Retrieval (GAR) for answering open-domain questions, which augments a query through text generation of heuristically discovered relevant contexts without external resources as supervision.

317, TITLE: Check It Again: Progressive Visual Question Answering via Visual Entailment

<https://aclanthology.org/2021.acl-long.317>

AUTHORS: Qingyi Si, Zheng Lin, Ming yu Zheng, Peng Fu, Weiping Wang  
HIGHLIGHT: In this paper, we propose a select-and-rerank (SAR) progressive framework based on Visual Entailment.

318, TITLE: A Mutual Information Maximization Approach for the Spurious Solution Problem in Weakly Supervised Question Answering

<https://aclanthology.org/2021.acl-long.318>

AUTHORS: Zhihong Shao, Lifeng Shang, Qun Liu, Minlie Huang  
HIGHLIGHT: In this paper, to alleviate the spurious solution problem, we propose to explicitly exploit such semantic correlations by maximizing the mutual information between question-answer pairs and predicted solutions.

319, TITLE: Breaking Down Walls of Text: How Can NLP Benefit Consumer Privacy?

<https://aclanthology.org/2021.acl-long.319>

AUTHORS: Abhilasha Ravichander, Alan W Black, Thomas Norton, Shomir Wilson, Norman Sadeh  
HIGHLIGHT: Our goal is to provide a roadmap for the development and use of language technologies to empower users to reclaim control over their privacy, limit privacy harms, and rally research efforts from the community towards addressing an issue with large social impact.

320, TITLE: Supporting Land Reuse of Former Open Pit Mining Sites using Text Classification and Active Learning

<https://aclanthology.org/2021.acl-long.320>

AUTHORS: Christopher Schr?der, Kim B?rgl, Yves Annanias, Andreas Niekler, Lydia M?ller, Daniel Wiegrefte, Christian Bender, Christoph Mengers, Gerik Scheuermann, Gerhard Heyer  
HIGHLIGHT: In this work we present and evaluate an automated workflow for supporting the post-mining management of former lignite open pit mines in the eastern part of Germany, where prior to any planned land reuse, aforementioned information has to be acquired to ensure the safety and validity of such an endeavor.

321, TITLE: Reliability Testing for Natural Language Processing Systems

<https://aclanthology.org/2021.acl-long.321>

AUTHORS: Samson Tan, Shafiq Joty, Kathy Baxter, Araz Taeihagh, Gregory A. Bennett, Min-Yen Kan  
HIGHLIGHT: To address this, we argue for the need for reliability testing and contextualize it among existing work on improving accountability.

322, TITLE: Learning Language and Multimodal Privacy-Preserving Markers of Mood from Mobile Data

<https://aclanthology.org/2021.acl-long.322>

AUTHORS: Paul Pu Liang, Terrance Liu, Anna Cai, Michal Muszynski, Ryo Ishii, Nick Allen, Randy Auerbach, David Brent, Ruslan Salakhutdinov, Louis-Philippe Morency

- HIGHLIGHT:** In this paper, we study behavioral markers of daily mood using a recent dataset of mobile behaviors from adolescent populations at high risk of suicidal behaviors.
- 323, **TITLE:** Anonymisation Models for Text Data: State of the art, Challenges and Future Directions  
<https://aclanthology.org/2021.acl-long.323>  
**AUTHORS:** Pierre Lison, Ildik? Pil?n, David Sanchez, Montserrat Batet, Lilja ?vrelid  
**HIGHLIGHT:** This position paper investigates the problem of automated text anonymisation, which is a prerequisite for secure sharing of documents containing sensitive information about individuals.
- 324, **TITLE:** End-to-End AMR Coreference Resolution  
<https://aclanthology.org/2021.acl-long.324>  
**AUTHORS:** Qiankun Fu, Linfeng Song, Wenyu Du, Yue Zhang  
**HIGHLIGHT:** We introduce the first end-to-end AMR coreference resolution model in order to build multi-sentence AMRs.
- 325, **TITLE:** How is BERT surprised? Layerwise detection of linguistic anomalies  
<https://aclanthology.org/2021.acl-long.325>  
**AUTHORS:** Bai Li, Zining Zhu, Guillaume Thomas, Yang Xu, Frank Rudzicz  
**HIGHLIGHT:** In this work, we use Gaussian models for density estimation at intermediate layers of three language models (BERT, RoBERTa, and XLNet), and evaluate our method on BLiMP, a grammaticality judgement benchmark.
- 326, **TITLE:** Psycholinguistic Tripartite Graph Network for Personality Detection  
<https://aclanthology.org/2021.acl-long.326>  
**AUTHORS:** Tao Yang, Feifan Yang, Haolan Ouyang, Xiaojun Quan  
**HIGHLIGHT:** In this paper, we propose a psycholinguistic knowledge-based tripartite graph network, TrigNet, which consists of a tripartite graph network and a BERT-based graph initializer.
- 327, **TITLE:** Verb Metaphor Detection via Contextual Relation Learning  
<https://aclanthology.org/2021.acl-long.327>  
**AUTHORS:** Wei Song, Shuhui Zhou, Ruiji Fu, Ting Liu, Lizhen Liu  
**HIGHLIGHT:** In this paper, we argue that it is appropriate to view this task as relation classification between a verb and its various contexts.
- 328, **TITLE:** Improving Speech Translation by Understanding and Learning from the Auxiliary Text Translation Task  
<https://aclanthology.org/2021.acl-long.328>  
**AUTHORS:** Yun Tang, Juan Pino, Xian Li, Changhan Wang, Dmitriy Genzel  
**HIGHLIGHT:** In this study, we are interested in training a speech translation model along with an auxiliary text translation task.
- 329, **TITLE:** Probing Toxic Content in Large Pre-Trained Language Models  
<https://aclanthology.org/2021.acl-long.329>  
**AUTHORS:** Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, Dit-Yan Yeung  
**HIGHLIGHT:** We propose a method based on logistic regression classifiers to probe English, French, and Arabic PTLMs and quantify the potentially harmful content that they convey with respect to a set of templates.
- 330, **TITLE:** Societal Biases in Language Generation: Progress and Challenges  
<https://aclanthology.org/2021.acl-long.330>  
**AUTHORS:** Emily Sheng, Kai-Wei Chang, Prem Natarajan, Nanyun Peng  
**HIGHLIGHT:** To better understand these challenges, we present a survey on societal biases in language generation, focusing on how data and techniques contribute to biases and progress towards reducing biases.
- 331, **TITLE:** Reservoir Transformers  
<https://aclanthology.org/2021.acl-long.331>  
**AUTHORS:** Sheng Shen, Alexei Baevski, Ari Morcos, Kurt Keutzer, Michael Auli, Douwe Kiela  
**HIGHLIGHT:** Inspired by old and well-established ideas in machine learning, we explore a variety of non-linear "reservoir" layers interspersed with regular transformer layers, and show improvements in wall-clock compute time until convergence, as well as overall performance, on various machine translation and (masked) language modelling tasks.
- 332, **TITLE:** Subsequence Based Deep Active Learning for Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.332>  
**AUTHORS:** Puria Radmard, Yassir Fathullah, Aldo Lipani

**HIGHLIGHT:** In this work, we overcome these limitations by allowing the AL algorithm to query subsequences within sentences, and propagate their labels to other sentences.

333, **TITLE:** Convolutions and Self-Attention: Re-interpreting Relative Positions in Pre-trained Language Models  
<https://aclanthology.org/2021.acl-long.333>  
**AUTHORS:** Tyler Chang, Yifan Xu, Weijian Xu, Zhuowen Tu  
**HIGHLIGHT:** In this paper, we detail the relationship between convolutions and self-attention in natural language tasks.

334, **TITLE:** BinaryBERT: Pushing the Limit of BERT Quantization  
<https://aclanthology.org/2021.acl-long.334>  
**AUTHORS:** Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jin Jin, Xin Jiang, Qun Liu, Michael Lyu, Irwin King  
**HIGHLIGHT:** In this paper, we propose BinaryBERT, which pushes BERT quantization to the limit by weight binarization.

335, **TITLE:** Are Pretrained Convolutions Better than Pretrained Transformers?  
<https://aclanthology.org/2021.acl-long.335>  
**AUTHORS:** Yi Tay, Mostafa Dehghani, Jai Prakash Gupta, Vamsi Aribandi, Dara Bahri, Zhen Qin, Donald Metzler  
**HIGHLIGHT:** Overall, the findings outlined in this paper suggest that conflating pre-training and architectural advances is misguided and that both advances should be considered independently.

336, **TITLE:** PairRE: Knowledge Graph Embeddings via Paired Relation Vectors  
<https://aclanthology.org/2021.acl-long.336>  
**AUTHORS:** Linlin Chao, Jianshan He, Taifeng Wang, Wei Chu  
**HIGHLIGHT:** To mitigate this problem, we propose PairRE, a model with paired vectors for each relation representation.

337, **TITLE:** Hierarchy-aware Label Semantics Matching Network for Hierarchical Text Classification  
<https://aclanthology.org/2021.acl-long.337>  
**AUTHORS:** Haibin Chen, Qianli Ma, Zhenxi Lin, Jianguye Yan  
**HIGHLIGHT:** To this end, we formulate the text-label semantics relationship as a semantic matching problem and thus propose a hierarchy-aware label semantics matching network (HiMatch).

338, **TITLE:** HiddenCut: Simple Data Augmentation for Natural Language Understanding with Better Generalizability  
<https://aclanthology.org/2021.acl-long.338>  
**AUTHORS:** Jiaao Chen, Dinghan Shen, Weizhu Chen, Diyi Yang  
**HIGHLIGHT:** To this end, we propose a simple yet effective data augmentation technique, HiddenCut, to better regularize the model and encourage it to learn more generalizable features.

339, **TITLE:** Neural Stylistic Response Generation with Disentangled Latent Variables  
<https://aclanthology.org/2021.acl-long.339>  
**AUTHORS:** Qingfu Zhu, Wei-Nan Zhang, Ting Liu, William Yang Wang  
**HIGHLIGHT:** In this paper, we propose to disentangle the content and style in latent space by diluting sentence-level information in style representations.

340, **TITLE:** Intent Classification and Slot Filling for Privacy Policies  
<https://aclanthology.org/2021.acl-long.340>  
**AUTHORS:** Wasi Ahmad, Jianfeng Chi, Tu Le, Thomas Norton, Yuan Tian, Kai-Wei Chang  
**HIGHLIGHT:** We present two alternative neural approaches as baselines, (1) intent classification and slot filling as a joint sequence tagging and (2) modeling them as a sequence-to-sequence (Seq2Seq) learning task.

341, **TITLE:** RADDLE: An Evaluation Benchmark and Analysis Platform for Robust Task-oriented Dialog Systems  
<https://aclanthology.org/2021.acl-long.341>  
**AUTHORS:** Baolin Peng, Chunyuan Li, Zhu Zhang, Chenguang Zhu, Jinchao Li, Jianfeng Gao  
**HIGHLIGHT:** In pursuit of these goals, we introduce the RADDLE benchmark, a collection of corpora and tools for evaluating the performance of models across a diverse set of domains.

342, **TITLE:** Semantic Representation for Dialogue Modeling  
<https://aclanthology.org/2021.acl-long.342>  
**AUTHORS:** Xuefeng Bai, Yulong Chen, Linfeng Song, Yue Zhang  
**HIGHLIGHT:** We develop an algorithm to construct dialogue-level AMR graphs from sentence-level AMRs and explore two ways to incorporate AMRs into dialogue systems.

343, TITLE: A Pre-training Strategy for Zero-Resource Response Selection in Knowledge-Grounded Conversations  
<https://aclanthology.org/2021.acl-long.343>  
AUTHORS: Chongyang Tao, Changyu Chen, Jiazhan Feng, Ji-Rong Wen, Rui Yan  
HIGHLIGHT: To overcome the challenge, we consider decomposing the training of the knowledge-grounded response selection into three tasks including: 1) query-passage matching task; 2) query-dialogue history matching task; 3) multi-turn response matching task, and joint learning all these tasks in a unified pre-trained language model.

344, TITLE: Dependency-driven Relation Extraction with Attentive Graph Convolutional Networks  
<https://aclanthology.org/2021.acl-long.344>  
AUTHORS: Yuanhe Tian, Guimin Chen, Yan Song, Xiang Wan  
HIGHLIGHT: In this paper, we propose a dependency-driven approach for relation extraction with attentive graph convolutional networks (A-GCN).

345, TITLE: Evaluating Entity Disambiguation and the Role of Popularity in Retrieval-Based NLP  
<https://aclanthology.org/2021.acl-long.345>  
AUTHORS: Anthony Chen, Pallavi Gudipati, Shayne Longpre, Xiao Ling, Sameer Singh  
HIGHLIGHT: We propose an evaluation benchmark for assessing the entity disambiguation capabilities of these retrievers, which we call Ambiguous Entity Retrieval (AmBER) sets.

346, TITLE: Evaluation Examples are not Equally Informative: How should that change NLP Leaderboards?  
<https://aclanthology.org/2021.acl-long.346>  
AUTHORS: Pedro Rodriguez, Joe Barrow, Alexander Miserlis Hoyle, John P. Lalor, Robin Jia, Jordan Boyd-Graber  
HIGHLIGHT: Building on educational testing, we create a Bayesian leaderboard model where latent subject skill and latent item difficulty predict correct responses.

347, TITLE: Claim Matching Beyond English to Scale Global Fact-Checking  
<https://aclanthology.org/2021.acl-long.347>  
AUTHORS: Ashkan Kazemi, Kiran Garimella, Devin Gaffney, Scott Hale  
HIGHLIGHT: In this paper, we discuss claim matching as a possible solution to scale fact-checking.

348, TITLE: SemFace: Pre-training Encoder and Decoder with a Semantic Interface for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.348>  
AUTHORS: Shuo Ren, Long Zhou, Shujie Liu, Furu Wei, Ming Zhou, Shuai Ma  
HIGHLIGHT: In this paper, we propose a better pre-training method for NMT by defining a semantic interface (SemFace) between the pre-trained encoder and the pre-trained decoder.

349, TITLE: Energy-Based Reranking: Improving Neural Machine Translation Using Energy-Based Models  
<https://aclanthology.org/2021.acl-long.349>  
AUTHORS: Sumanta Bhattacharyya, Amirmohammad Rooshenas, Subhajit Naskar, Simeng Sun, Mohit Iyyer, Andrew McCallum  
HIGHLIGHT: To benefit from this observation, we train an energy-based model to mimic the behavior of the task measure (i.e., the energy-based model assigns lower energy to samples with higher BLEU score), which is resulted in a re-ranking algorithm based on the samples drawn from NMT: energy-based re-ranking (EBR).

350, TITLE: Syntax-augmented Multilingual BERT for Cross-lingual Transfer  
<https://aclanthology.org/2021.acl-long.350>  
AUTHORS: Wasi Ahmad, Haoran Li, Kai-Wei Chang, Yashar Mehdad  
HIGHLIGHT: This work shows that explicitly providing language syntax and training mBERT using an auxiliary objective to encode the universal dependency tree structure helps cross-lingual transfer.

351, TITLE: How to Adapt Your Pretrained Multilingual Model to 1600 Languages  
<https://aclanthology.org/2021.acl-long.351>  
AUTHORS: Abteen Ebrahimi, Katharina Kann  
HIGHLIGHT: In this paper, we evaluate the performance of existing methods to adapt PMMs to new languages using a resource available for close to 1600 languages: the New Testament.

352, TITLE: Weakly Supervised Named Entity Tagging with Learnable Logical Rules  
<https://aclanthology.org/2021.acl-long.352>  
AUTHORS: Jiacheng Li, Haibo Ding, Jingbo Shang, Julian McAuley, Zhe Feng

**HIGHLIGHT:** In this work, we propose a novel method TALLOR that bootstraps high-quality logical rules to train a neural tagger in a fully automated manner.

353, **TITLE:** Prefix-Tuning: Optimizing Continuous Prompts for Generation

<https://aclanthology.org/2021.acl-long.353>

**AUTHORS:** Xiang Lisa Li, Percy Liang

**HIGHLIGHT:** In this paper, we propose prefix-tuning, a lightweight alternative to fine-tuning for natural language generation tasks, which keeps language model parameters frozen and instead optimizes a sequence of continuous task-specific vectors, which we call the prefix.

354, **TITLE:** One2Set: Generating Diverse Keyphrases as a Set

<https://aclanthology.org/2021.acl-long.354>

**AUTHORS:** Jiacheng Ye, Tao Gui, Yichao Luo, Yige Xu, Qi Zhang

**HIGHLIGHT:** In this work, we propose a new training paradigm One2Set without predefining an order to concatenate the keyphrases.

355, **TITLE:** Continuous Language Generative Flow

<https://aclanthology.org/2021.acl-long.355>

**AUTHORS:** Zineng Tang, Shiyue Zhang, Hyounghun Kim, Mohit Bansal

**HIGHLIGHT:** In this paper, we propose a flow-based language generation model by adapting previous flow generative models to language generation via continuous input embeddings, adapted affine coupling structures, and a novel architecture for autoregressive text generation.

356, **TITLE:** TWAG: A Topic-Guided Wikipedia Abstract Generator

<https://aclanthology.org/2021.acl-long.356>

**AUTHORS:** Fangwei Zhu, Shangqing Tu, Jiaxin Shi, Juanzi Li, Lei Hou, Tong Cui

**HIGHLIGHT:** In this paper, we propose a two-stage model TWAG that guides the abstract generation with topical information.

357, **TITLE:** ForecastQA: A Question Answering Challenge for Event Forecasting with Temporal Text Data

<https://aclanthology.org/2021.acl-long.357>

**AUTHORS:** Woojeong Jin, Rahul Khanna, Suji Kim, Dong-Ho Lee, Fred Morstatter, Aram Galstyan, Xiang Ren

**HIGHLIGHT:** To showcase the usefulness of this task formulation, we introduce ForecastQA, a question-answering dataset consisting of 10,392 event forecasting questions, which have been collected and verified via crowdsourcing efforts.

358, **TITLE:** Recursive Tree-Structured Self-Attention for Answer Sentence Selection

<https://aclanthology.org/2021.acl-long.358>

**AUTHORS:** Khalil Mrini, Emilia Farcas, Ndapa Nakashole

**HIGHLIGHT:** We introduce the Tree Aggregation Transformer: a novel recursive, tree-structured self-attention model for AS2.

359, **TITLE:** How Knowledge Graph and Attention Help? A Qualitative Analysis into Bag-level Relation Extraction

<https://aclanthology.org/2021.acl-long.359>

**AUTHORS:** Zikun Hu, Yixin Cao, Lifu Huang, Tat-Seng Chua

**HIGHLIGHT:** In this paper, we contribute a dataset and propose a paradigm to quantitatively evaluate the effect of attention and KG on bag-level relation extraction (RE).

360, **TITLE:** Trigger is Not Sufficient: Exploiting Frame-aware Knowledge for Implicit Event Argument Extraction

<https://aclanthology.org/2021.acl-long.360>

**AUTHORS:** Kaiwen Wei, Xian Sun, Zequn Zhang, Jingyuan Zhang, Guo Zhi, Li Jin

**HIGHLIGHT:** In this work, we present a Frame-aware Event Argument Extraction (FEAE) learning framework to tackle this issue through reasoning in event frame-level scope.

361, **TITLE:** Element Intervention for Open Relation Extraction

<https://aclanthology.org/2021.acl-long.361>

**AUTHORS:** Fangchao Liu, Lingyong Yan, Hongyu Lin, Xianpei Han, Le Sun

**HIGHLIGHT:** In this paper, we revisit the procedure of OpenRE from a causal view.

362, **TITLE:** AdaTag: Multi-Attribute Value Extraction from Product Profiles with Adaptive Decoding

<https://aclanthology.org/2021.acl-long.362>

AUTHORS: Jun Yan, Nasser Zalmout, Yan Liang, Christan Grant, Xiang Ren, Xin Luna Dong  
HIGHLIGHT: In this paper we present AdaTag, which uses adaptive decoding to handle extraction.

363, TITLE: CoRI: Collective Relation Integration with Data Augmentation for Open Information Extraction  
<https://aclanthology.org/2021.acl-long.363>  
AUTHORS: Zhengbao Jiang, Jialong Han, Bunyamin Sisman, Xin Luna Dong  
HIGHLIGHT: We propose a two-stage Collective Relation Integration (CoRI) model, where the first stage independently makes candidate predictions, and the second stage employs a collective model that accesses all candidate predictions to make globally coherent predictions.

364, TITLE: Benchmarking Scalable Methods for Streaming Cross Document Entity Coreference  
<https://aclanthology.org/2021.acl-long.364>  
AUTHORS: Robert L Logan IV, Andrew McCallum, Sameer Singh, Dan Bikel  
HIGHLIGHT: In this work, we address this issue by compiling a large benchmark adapted from existing free datasets, and performing a comprehensive evaluation of a number of novel and existing baseline models.

365, TITLE: Search from History and Reason for Future: Two-stage Reasoning on Temporal Knowledge Graphs  
<https://aclanthology.org/2021.acl-long.365>  
AUTHORS: Zixuan Li, Xiaolong Jin, Saiping Guan, Wei Li, Jiafeng Guo, Yuanzhuo Wang, Xueqi Cheng  
HIGHLIGHT: Inspired by this mechanism, we propose CluSTeR to predict future facts in a two-stage manner, Clue Searching and Temporal Reasoning, accordingly.

366, TITLE: Employing Argumentation Knowledge Graphs for Neural Argument Generation  
<https://aclanthology.org/2021.acl-long.366>  
AUTHORS: Khalid Al Khatib, Lukas Trautner, Henning Wachsmuth, Yufang Hou, Benno Stein  
HIGHLIGHT: Motivated by the effectiveness of utilizing knowledge graphs for supporting general text generation tasks, this paper investigates the usage of argumentation-related knowledge graphs to control the generation of arguments.

367, TITLE: Learning Span-Level Interactions for Aspect Sentiment Triplet Extraction  
<https://aclanthology.org/2021.acl-long.367>  
AUTHORS: Lu Xu, Yew Ken Chia, Lidong Bing  
HIGHLIGHT: To ease the high computational cost caused by span enumeration, we propose a dual-channel span pruning strategy by incorporating supervision from the Aspect Term Extraction (ATE) and Opinion Term Extraction (OTE) tasks.

368, TITLE: On Compositional Generalization of Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.368>  
AUTHORS: Yafu Li, Yongjing Yin, Yulong Chen, Yue Zhang  
HIGHLIGHT: In this paper, we study NMT models from the perspective of compositional generalization by building a benchmark dataset, CoGnition, consisting of 216k clean and consistent sentence pairs.

369, TITLE: Mask-Align: Self-Supervised Neural Word Alignment  
<https://aclanthology.org/2021.acl-long.369>  
AUTHORS: Chi Chen, Maosong Sun, Yang Liu  
HIGHLIGHT: In this paper, we propose Mask-Align, a self-supervised word alignment model that takes advantage of the full context on the target side.

370, TITLE: GWLAN: General Word-Level AutocompletiON for Computer-Aided Translation  
<https://aclanthology.org/2021.acl-long.370>  
AUTHORS: Huayang Li, Lemao Liu, Guoping Huang, Shuming Shi  
HIGHLIGHT: In this paper, we propose the task of general word-level autocompletion (GWLAN) from a real-world CAT scenario, and construct the first public benchmark to facilitate research in this topic.

371, TITLE: De-biasing Distantly Supervised Named Entity Recognition via Causal Intervention  
<https://aclanthology.org/2021.acl-long.371>  
AUTHORS: Wenkai Zhang, Hongyu Lin, Xianpei Han, Le Sun  
HIGHLIGHT: In this paper, we fundamentally explain the dictionary bias via a Structural Causal Model (SCM), categorize the bias into intra-dictionary and inter-dictionary biases, and identify their causes.

372, TITLE: A Span-Based Model for Joint Overlapped and Discontinuous Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.372>

- AUTHORS: Fei Li, ZhiChao Lin, Meishan Zhang, Donghong Ji  
HIGHLIGHT: In this paper, we propose a novel span-based model that can recognize both overlapped and discontinuous entities jointly.
- 373, TITLE: MLBiNet: A Cross-Sentence Collective Event Detection Network  
<https://aclanthology.org/2021.acl-long.373>  
AUTHORS: Dongfang Lou, Zhilin Liao, Shumin Deng, Ningyu Zhang, Huajun Chen  
HIGHLIGHT: In this paper, we reformulate it as a Seq2Seq task and propose a Multi-Layer Bidirectional Network (MLBiNet) to capture the document-level association of events and semantic information simultaneously.
- 374, TITLE: Exploiting Document Structures and Cluster Consistencies for Event Coreference Resolution  
<https://aclanthology.org/2021.acl-long.374>  
AUTHORS: Hieu Minh Tran, Duy Phung, Thien Huu Nguyen  
HIGHLIGHT: This work addresses such limitations by introducing a novel deep learning model for ECR.
- 375, TITLE: StereoRel: Relational Triple Extraction from a Stereoscopic Perspective  
<https://aclanthology.org/2021.acl-long.375>  
AUTHORS: Xuetao Tian, Liping Jing, Lu He, Feng Liu  
HIGHLIGHT: To intuitively explore the above issues and address them, in this paper, we provide a revealing insight into relational triple extraction from a stereoscopic perspective, which rationalizes the occurrence of these issues and exposes the shortcomings of existing methods.
- 376, TITLE: Knowledge-Enriched Event Causality Identification via Latent Structure Induction Networks  
<https://aclanthology.org/2021.acl-long.376>  
AUTHORS: Pengfei Cao, Xinyu Zuo, Yubo Chen, Kang Liu, Jun Zhao, Yuguang Chen, Weihua Peng  
HIGHLIGHT: Inspired by it, we propose a novel Latent Structure Induction Network (LSIN) to incorporate the external structural knowledge into this task.
- 377, TITLE: Turn the Combination Lock: Learnable Textual Backdoor Attacks via Word Substitution  
<https://aclanthology.org/2021.acl-long.377>  
AUTHORS: Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, Maosong Sun  
HIGHLIGHT: In this work, we present invisible backdoors that are activated by a learnable combination of word substitution.
- 378, TITLE: Parameter-Efficient Transfer Learning with Diff Pruning  
<https://aclanthology.org/2021.acl-long.378>  
AUTHORS: Demi Guo, Alexander Rush, Yoon Kim  
HIGHLIGHT: The approach learns a task-specific "diff" vector that extends the original pretrained parameters.
- 379, TITLE: R2D2: Recursive Transformer based on Differentiable Tree for Interpretable Hierarchical Language Modeling  
<https://aclanthology.org/2021.acl-long.379>  
AUTHORS: Xiang Hu, Haitao Mi, Zujie Wen, Yafang Wang, Yi Su, Jing Zheng, Gerard de Melo  
HIGHLIGHT: In this paper, we propose a recursive Transformer model based on differentiable CKY style binary trees to emulate this composition process, and we extend the bidirectional language model pre-training objective to this architecture, attempting to predict each word given its left and right abstraction nodes.
- 380, TITLE: Risk Minimization for Zero-shot Sequence Labeling  
<https://aclanthology.org/2021.acl-long.380>  
AUTHORS: Zechuan Hu, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, Kewei Tu  
HIGHLIGHT: In this paper, we propose a novel unified framework for zero-shot sequence labeling with minimum risk training and design a new decomposable risk function that models the relations between the predicted labels from the source models and the true labels.
- 381, TITLE: WARP: Word-level Adversarial ReProgramming  
<https://aclanthology.org/2021.acl-long.381>  
AUTHORS: Karen Hambardzumyan, Hrant Khachatrian, Jonathan May  
HIGHLIGHT: In this paper, we present an alternative approach based on adversarial reprogramming, which extends earlier work on automatic prompt generation.
- 382, TITLE: Lexicon Learning for Few Shot Sequence Modeling  
<https://aclanthology.org/2021.acl-long.382>

AUTHORS: Ekin Akyurek, Jacob Andreas  
HIGHLIGHT: To address this, we augment neural decoders with a lexical translation mechanism that generalizes existing copy mechanisms to incorporate learned, decontextualized, token-level translation rules.

383, TITLE: Personalized Transformer for Explainable Recommendation  
<https://aclanthology.org/2021.acl-long.383>  
AUTHORS: Lei Li, Yongfeng Zhang, Li Chen  
HIGHLIGHT: To address this problem, we present a PErsonalized Transformer for Explainable Recommendation (PETER), on which we design a simple and effective learning objective that utilizes the IDs to predict the words in the target explanation, so as to endow the IDs with linguistic meanings and to achieve personalized Transformer.

384, TITLE: Generating SOAP Notes from Doctor-Patient Conversations Using Modular Summarization Techniques  
<https://aclanthology.org/2021.acl-long.384>  
AUTHORS: Kundan Krishna, Sopan Khosla, Jeffrey Bigham, Zachary C. Lipton  
HIGHLIGHT: In this paper, we introduce the first complete pipelines to leverage deep summarization models to generate these notes based on transcripts of conversations between physicians and patients.

385, TITLE: Tail-to-Tail Non-Autoregressive Sequence Prediction for Chinese Grammatical Error Correction  
<https://aclanthology.org/2021.acl-long.385>  
AUTHORS: Piji Li, Shuming Shi  
HIGHLIGHT: We investigate the problem of Chinese Grammatical Error Correction (CGEC) and present a new framework named Tail-to-Tail (TtT) non-autoregressive sequence prediction to address the deep issues hidden in CGEC.

386, TITLE: Early Detection of Sexual Predators in Chats  
<https://aclanthology.org/2021.acl-long.386>  
AUTHORS: Matthias Vogt, Ulf Leser, Alan Akbik  
HIGHLIGHT: In this work, we instead investigate this problem from the point of view of prevention.

387, TITLE: Writing by Memorizing: Hierarchical Retrieval-based Medical Report Generation  
<https://aclanthology.org/2021.acl-long.387>  
AUTHORS: Xingyi Yang, Muchao Ye, Quanzeng You, Fenglong Ma  
HIGHLIGHT: To address these issues, we propose MedWriter that incorporates a novel hierarchical retrieval mechanism to automatically extract both report and sentence-level templates for clinically accurate report generation.

388, TITLE: Concept-Based Label Embedding via Dynamic Routing for Hierarchical Text Classification  
<https://aclanthology.org/2021.acl-long.388>  
AUTHORS: Xuepeng Wang, Li Zhao, Bing Liu, Tao Chen, Feng Zhang, Di Wang  
HIGHLIGHT: In this paper, we propose a novel concept-based label embedding method that can explicitly represent the concept and model the sharing mechanism among classes for the hierarchical text classification.

389, TITLE: VisualSparta: An Embarrassingly Simple Approach to Large-scale Text-to-Image Search with Weighted Bag-of-words  
<https://aclanthology.org/2021.acl-long.389>  
AUTHORS: Xiaopeng Lu, Tiancheng Zhao, Kyusong Lee  
HIGHLIGHT: In this paper, we propose VisualSparta, a novel (Visual-text Sparse Transformer Matching) model that shows significant improvement in terms of both accuracy and efficiency.

390, TITLE: Few-Shot Text Ranking with Meta Adapted Synthetic Weak Supervision  
<https://aclanthology.org/2021.acl-long.390>  
AUTHORS: Si Sun, Yingzhuo Qian, Zhenghao Liu, Chenyan Xiong, Kaitao Zhang, Jie Bao, Zhiyuan Liu, Paul Bennett  
HIGHLIGHT: To democratize the benefits of Neu-IR, this paper presents MetaAdaptRank, a domain adaptive learning method that generalizes Neu-IR models from label-rich source domains to few-shot target domains.

391, TITLE: Semi-Supervised Text Classification with Balanced Deep Representation Distributions  
<https://aclanthology.org/2021.acl-long.391>  
AUTHORS: Changchun Li, Ximing Li, Jihong Ouyang  
HIGHLIGHT: With this insight, we propose a novel SSTC method, namely Semi-Supervised Text Classification with Balanced Deep representation Distributions (S2TC-BDD).

392, TITLE: Improving Document Representations by Generating Pseudo Query Embeddings for Dense Retrieval



<https://aclanthology.org/2021.acl-long.392>

AUTHORS: Hongyin Tang, Xingwu Sun, Beihong Jin, Jingang Wang, Fuzheng Zhang, Wei Wu  
HIGHLIGHT: To address this problem, we design a method to mimic the queries to each of the documents by an iterative clustering process and represent the documents by multiple pseudo queries (i.e., the cluster centroids).

393, TITLE: ConSERT: A Contrastive Framework for Self-Supervised Sentence Representation Transfer

<https://aclanthology.org/2021.acl-long.393>

AUTHORS: Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, Weiran Xu  
HIGHLIGHT: In this paper, we present ConSERT, a Contrastive Framework for Self-Supervised SEntence Representation Transfer, that adopts contrastive learning to fine-tune BERT in an unsupervised and effective way.

394, TITLE: Exploring Dynamic Selection of Branch Expansion Orders for Code Generation

<https://aclanthology.org/2021.acl-long.394>

AUTHORS: Hui Jiang, Chulun Zhou, Fandong Meng, Biao Zhang, Jie Zhou, Degen Huang, Qingqiang Wu, Jinsong Su  
HIGHLIGHT: In this paper, we propose to equip the Seq2Tree model with a context-based Branch Selector, which is able to dynamically determine optimal expansion orders of branches for multi-branch nodes.

395, TITLE: COINS: Dynamically Generating COntextualized Inference Rules for Narrative Story Completion

<https://aclanthology.org/2021.acl-long.395>

AUTHORS: Debjit Paul, Anette Frank  
HIGHLIGHT: In this paper we present Coins, a recursive inference framework that i) iteratively reads context sentences, ii) dynamically generates contextualized inference rules, encodes them, and iii) uses them to guide task-specific output generation.

396, TITLE: Reasoning over Entity-Action-Location Graph for Procedural Text Understanding

<https://aclanthology.org/2021.acl-long.396>

AUTHORS: Hao Huang, Xiubo Geng, Jian Pei, Guodong Long, Daxin Jiang  
HIGHLIGHT: In this paper, we propose a novel approach (REAL) to procedural text understanding, where we build a general framework to systematically model the entity-entity, entity-action, and entity-location relations using a graph neural network.

397, TITLE: From Paraphrasing to Semantic Parsing: Unsupervised Semantic Parsing via Synchronous Semantic Decoding

<https://aclanthology.org/2021.acl-long.397>

AUTHORS: Shan Wu, Bo Chen, Chunlei Xin, Xianpei Han, Le Sun, Weipeng Zhang, Jiansong Chen, Fan Yang, Xunliang Cai  
HIGHLIGHT: In this paper, we propose an unsupervised semantic parsing method - Synchronous Semantic Decoding (SSD), which can simultaneously resolve the semantic gap and the structure gap by jointly leveraging paraphrasing and grammar-constrained decoding.

398, TITLE: Pre-training Universal Language Representation

<https://aclanthology.org/2021.acl-long.398>

AUTHORS: Yian Li, Hai Zhao  
HIGHLIGHT: This work introduces universal language representation learning, i.e., embeddings of different levels of linguistic units or text with quite diverse lengths in a uniform vector space.

399, TITLE: Structural Pre-training for Dialogue Comprehension

<https://aclanthology.org/2021.acl-long.399>

AUTHORS: Zhuosheng Zhang, Hai Zhao  
HIGHLIGHT: In this work, we present SPIDER, Structural Pre-trained DialoguE Reader, to capture dialogue exclusive features.

400, TITLE: AutoTinyBERT: Automatic Hyper-parameter Optimization for Efficient Pre-trained Language Models

<https://aclanthology.org/2021.acl-long.400>

AUTHORS: Yichun Yin, Cheng Chen, Lifeng Shang, Xin Jiang, Xiao Chen, Qun Liu  
HIGHLIGHT: In this paper, we adopt the one-shot Neural Architecture Search (NAS) to automatically search architecture hyper-parameters.

401, TITLE: Data Augmentation with Adversarial Training for Cross-Lingual NLI

<https://aclanthology.org/2021.acl-long.401>

AUTHORS: Xin Dong, Yaxin Zhu, Zuohui Fu, Dongkuan Xu, Gerard de Melo  
HIGHLIGHT: In this paper, we propose a novel data augmentation strategy for better cross-lingual natural language inference by enriching the data to reflect more diversity in a semantically faithful way.

- 402, TITLE: Bootstrapped Unsupervised Sentence Representation Learning  
<https://aclanthology.org/2021.acl-long.402>  
AUTHORS: Yan Zhang, Ruidan He, Zuozhu Liu, Lidong Bing, Haizhou Li  
HIGHLIGHT: In this paper, we propose a new framework with a two-branch Siamese Network which maximizes the similarity between two augmented views of each sentence.
- 403, TITLE: Learning Event Graph Knowledge for Abductive Reasoning  
<https://aclanthology.org/2021.acl-long.403>  
AUTHORS: Li Du, Xiao Ding, Ting Liu, Bing Qin  
HIGHLIGHT: To fill this gap, we propose a variational autoencoder based model ege-RoBERTa, which employs a latent variable to capture the necessary commonsense knowledge from event graph for guiding the abductive reasoning task.
- 404, TITLE: A Cognitive Regularizer for Language Modeling  
<https://aclanthology.org/2021.acl-long.404>  
AUTHORS: Jason Wei, Clara Meister, Ryan Cotterell  
HIGHLIGHT: In this work, we explore whether the UID hypothesis can be operationalized as an inductive bias for statistical language modeling.
- 405, TITLE: Lower Perplexity is Not Always Human-Like  
<https://aclanthology.org/2021.acl-long.405>  
AUTHORS: Tatsuki Kuribayashi, Yohei Oseki, Takumi Ito, Ryo Yoshida, Masayuki Asahara, Kentaro Inui  
HIGHLIGHT: In order to fill the gap, this paper investigates whether the established results in computational psycholinguistics can be generalized across languages.
- 406, TITLE: Word Sense Disambiguation: Towards Interactive Context Exploitation from Both Word and Sense Perspectives  
<https://aclanthology.org/2021.acl-long.406>  
AUTHORS: Ming Wang, Yinglin Wang  
HIGHLIGHT: In this paper, we convert the nearly isolated decisions into interrelated ones by exposing senses in context when learning sense embeddings in a similarity-based Sense Aware Context Exploitation (SACE) architecture.
- 407, TITLE: A Knowledge-Guided Framework for Frame Identification  
<https://aclanthology.org/2021.acl-long.407>  
AUTHORS: Xuefeng Su, Ru Li, Xiaoli Li, Jeff Z. Pan, Hu Zhang, Qinghua Chai, Xiaoqi Han  
HIGHLIGHT: In this paper, we propose a Knowledge-Guided Frame Identification framework (KGFI) that integrates three types frame knowledge, including frame definitions, frame elements and frame-to-frame relations, to learn better frame representation, which guides the KGFI to jointly map target words and frames into the same embedding space and subsequently identify the best frame by calculating the dot-product similarity scores between the target word embedding and all of the frame embeddings.
- 408, TITLE: Obtaining Better Static Word Embeddings Using Contextual Embedding Models  
<https://aclanthology.org/2021.acl-long.408>  
AUTHORS: Prakhar Gupta, Martin Jaggi  
HIGHLIGHT: In this work, we demonstrate that our proposed distillation method, which is a simple extension of CBOW-based training, allows to significantly improve computational efficiency of NLP applications, while outperforming the quality of existing static embeddings trained from scratch as well as those distilled from previously proposed methods.
- 409, TITLE: Meta-Learning with Variational Semantic Memory for Word Sense Disambiguation  
<https://aclanthology.org/2021.acl-long.409>  
AUTHORS: Yingjun Du, Nithin Holla, Xiantong Zhen, Cees Snoek, Ekaterina Shutova  
HIGHLIGHT: Aiming to further close this gap, we propose a model of semantic memory for WSD in a meta-learning setting.
- 410, TITLE: LexFit: Lexical Fine-Tuning of Pretrained Language Models  
<https://aclanthology.org/2021.acl-long.410>  
AUTHORS: Ivan Vulic, Edoardo Maria Ponti, Anna Korhonen, Goran Glava?  
HIGHLIGHT: Inspired by prior work on semantic specialization of static word embedding (WE) models, we show that it is possible to expose and enrich lexical knowledge from the LMs, that is, to specialize them to serve as effective and universal "decontextualized" word encoders even when fed input words "in isolation" (i.e., without any context).
- 411, TITLE: Text-Free Image-to-Speech Synthesis Using Learned Segmental Units  
<https://aclanthology.org/2021.acl-long.411>

AUTHORS: Wei-Ning Hsu, David Harwath, Tyler Miller, Christopher Song, James Glass  
HIGHLIGHT: In this paper we present the first model for directly synthesizing fluent, natural-sounding spoken audio captions for images that does not require natural language text as an intermediate representation or source of supervision.

412, TITLE: CTFN: Hierarchical Learning for Multimodal Sentiment Analysis Using Coupled-Translation Fusion Network  
<https://aclanthology.org/2021.acl-long.412>  
AUTHORS: Jiajia Tang, Kang Li, Xuanyu Jin, Andrzej Cichocki, Qibin Zhao, Wanzeng Kong  
HIGHLIGHT: In this work, the coupled-translation fusion network (CTFN) is firstly proposed to model bi-direction interplay via couple learning, ensuring the robustness in respect to missing modalities.

413, TITLE: Positional Artefacts Propagate Through Masked Language Model Embeddings  
<https://aclanthology.org/2021.acl-long.413>  
AUTHORS: Ziyang Luo, Artur Kulmizev, Xiaoxi Mao  
HIGHLIGHT: In this work, we demonstrate that the contextualized word vectors derived from pretrained masked language model-based encoders share a common, perhaps undesirable pattern across layers.

414, TITLE: Language Model Evaluation Beyond Perplexity  
<https://aclanthology.org/2021.acl-long.414>  
AUTHORS: Clara Meister, Ryan Cotterell  
HIGHLIGHT: We propose an alternate approach to quantifying how well language models learn natural language: we ask how well they match the statistical tendencies of natural language.

415, TITLE: Learning to Explain: Generating Stable Explanations Fast  
<https://aclanthology.org/2021.acl-long.415>  
AUTHORS: Xuelin Situ, Ingrid Zukerman, Cecile Paris, Sameen Maruf, Gholamreza Haffari  
HIGHLIGHT: In this paper, we propose a Learning to Explain (L2E) approach that learns the behaviour of an underlying explanation algorithm simultaneously from all training examples.

416, TITLE: StereoSet: Measuring stereotypical bias in pretrained language models  
<https://aclanthology.org/2021.acl-long.416>  
AUTHORS: Moin Nadeem, Anna Bethke, Siva Reddy  
HIGHLIGHT: We present StereoSet, a large-scale natural English dataset to measure stereotypical biases in four domains: gender, profession, race, and religion.

417, TITLE: Alignment Rationale for Natural Language Inference  
<https://aclanthology.org/2021.acl-long.417>  
AUTHORS: Zhongtao Jiang, Yuanzhe Zhang, Zhao Yang, Jun Zhao, Kang Liu  
HIGHLIGHT: To this end, this paper presents AREC, a post-hoc approach to generate alignment rationale explanations for co-attention based models in NLI.

418, TITLE: Enabling Lightweight Fine-tuning for Pre-trained Language Model Compression based on Matrix Product Operators  
<https://aclanthology.org/2021.acl-long.418>  
AUTHORS: Peiyu Liu, Ze-Feng Gao, Wayne Xin Zhao, Zhi-Yuan Xie, Zhong-Yi Lu, Ji-Rong Wen  
HIGHLIGHT: This paper presents a novel pre-trained language models (PLM) compression approach based on the matrix product operator (short as MPO) from quantum many-body physics.

419, TITLE: On Sample Based Explanation Methods for NLP: Faithfulness, Efficiency and Semantic Evaluation  
<https://aclanthology.org/2021.acl-long.419>  
AUTHORS: Wei Zhang, Ziming Huang, Yada Zhu, Guangnan Ye, Xiaodong Cui, Fan Zhang  
HIGHLIGHT: In this work, for the first time, we can improve the interpretability of explanations by allowing arbitrary text sequences as the explanation unit.

420, TITLE: Syntax-Enhanced Pre-trained Model  
<https://aclanthology.org/2021.acl-long.420>  
AUTHORS: Zenan Xu, Daya Guo, Duyu Tang, Qinliang Su, Linjun Shou, Ming Gong, Wanjun Zhong, Xiaojun Quan, Daxin Jiang, Nan Duan  
HIGHLIGHT: To address this, we present a model that utilizes the syntax of text in both pre-training and fine-tuning stages.

- 421, TITLE: Matching Distributions between Model and Data: Cross-domain Knowledge Distillation for Unsupervised Domain Adaptation  
<https://aclanthology.org/2021.acl-long.421>  
AUTHORS: Bo Zhang, Xiaoming Zhang, Yun Liu, Lei Cheng, Zhoujun Li  
HIGHLIGHT: We propose a generic framework named Cross-domain Knowledge Distillation (CdKD) without needing any source data.
- 422, TITLE: Counterfactual Inference for Text Classification Debiasing  
<https://aclanthology.org/2021.acl-long.422>  
AUTHORS: Chen Qian, Fuli Feng, Lijie Wen, Chunping Ma, Pengjun Xie  
HIGHLIGHT: Inspired by this, we propose a model-agnostic text classification debiasing framework - Corsair, which can effectively avoid employing data manipulations or designing balancing mechanisms.
- 423, TITLE: HieRec: Hierarchical User Interest Modeling for Personalized News Recommendation  
<https://aclanthology.org/2021.acl-long.423>  
AUTHORS: Tao Qi, Fangzhao Wu, Chuhan Wu, Peiru Yang, Yang Yu, Xing Xie, Yongfeng Huang  
HIGHLIGHT: In this paper, we propose a news recommendation method with hierarchical user interest modeling, named HieRec.
- 424, TITLE: PP-Rec: News Recommendation with Personalized User Interest and Time-aware News Popularity  
<https://aclanthology.org/2021.acl-long.424>  
AUTHORS: Tao Qi, Fangzhao Wu, Chuhan Wu, Yongfeng Huang  
HIGHLIGHT: Thus, in this paper we propose to incorporate news popularity information to alleviate the cold-start and diversity problems for personalized news recommendation.
- 425, TITLE: Article Reranking by Memory-Enhanced Key Sentence Matching for Detecting Previously Fact-Checked Claims  
<https://aclanthology.org/2021.acl-long.425>  
AUTHORS: Qiang Sheng, Juan Cao, Xueyao Zhang, Xirong Li, Lei Zhong  
HIGHLIGHT: In this paper, we propose a novel reranker, MTM (Memory-enhanced Transformers for Matching) to rank FC-articles using key sentences selected with event (lexical and semantic) and pattern information.
- 426, TITLE: Defense against Synonym Substitution-based Adversarial Attacks via Dirichlet Neighborhood Ensemble  
<https://aclanthology.org/2021.acl-long.426>  
AUTHORS: Yi Zhou, Xiaoqing Zheng, Cho-Jui Hsieh, Kai-Wei Chang, Xuanjing Huang  
HIGHLIGHT: We propose Dirichlet Neighborhood Ensemble (DNE), a randomized method for training a robust model to defense synonym substitution-based attacks.
- 427, TITLE: Shortformer: Better Language Modeling using Shorter Inputs  
<https://aclanthology.org/2021.acl-long.427>  
AUTHORS: Ofir Press, Noah A. Smith, Mike Lewis  
HIGHLIGHT: We identify conditions where shorter inputs are not harmful, and achieve perplexity and efficiency improvements through two new methods that decrease input length.
- 428, TITLE: BanditMTL: Bandit-based Multi-task Learning for Text Classification  
<https://aclanthology.org/2021.acl-long.428>  
AUTHORS: Yuren Mao, Zekai Wang, Weiwei Liu, Xuemin Lin, Wenbin Hu  
HIGHLIGHT: Accordingly, to fill this gap, this paper investigates how the task might be effectively regularized, and consequently proposes a multi-task learning method based on adversarial multi-armed bandit.
- 429, TITLE: Unified Interpretation of Softmax Cross-Entropy and Negative Sampling: With Case Study for Knowledge Graph Embedding  
<https://aclanthology.org/2021.acl-long.429>  
AUTHORS: Hidetaka Kamigaito, Katsuhiko Hayashi  
HIGHLIGHT: We attempted to solve this problem by using the Bregman divergence to provide a unified interpretation of the softmax cross-entropy and negative sampling loss functions.
- 430, TITLE: De-Confounded Variational Encoder-Decoder for Logical Table-to-Text Generation  
<https://aclanthology.org/2021.acl-long.430>  
AUTHORS: Wenqing Chen, Jidong Tian, Yitian Li, Hao He, Yaohui Jin

**HIGHLIGHT:** In this paper, we propose a de-confounded variational encoder-decoder (DCVED) based on causal intervention, learning the objective  $p(y|\text{do}(x))$ .

431, **TITLE:** Rethinking Stealthiness of Backdoor Attack against NLP Models

<https://aclanthology.org/2021.acl-long.431>

**AUTHORS:** Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, Xu Sun

**HIGHLIGHT:** In this work, we point out a potential problem of current backdoor attacking research: its evaluation ignores the stealthiness of backdoor attacks, and most of existing backdoor attacking methods are not stealthy either to system deployers or to system users.

432, **TITLE:** Crowdsourcing Learning as Domain Adaptation: A Case Study on Named Entity Recognition

<https://aclanthology.org/2021.acl-long.432>

**AUTHORS:** Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, Pengjun Xie

**HIGHLIGHT:** We take a different point in this work, regarding all crowdsourced annotations as gold-standard with respect to the individual annotators.

433, **TITLE:** Exploring Distantly-Labeled Rationales in Neural Network Models

<https://aclanthology.org/2021.acl-long.433>

**AUTHORS:** Quzhe Huang, Shengqi Zhu, Yansong Feng, Dongyan Zhao

**HIGHLIGHT:** In this paper, we propose two novel auxiliary loss functions to make better use of distantly-labeled rationales, which encourage models to maintain their focus on important words beyond labeled rationales (PINs) and alleviate redundant training on non-helpful rationales (NoIRs).

434, **TITLE:** Learning to Perturb Word Embeddings for Out-of-distribution QA

<https://aclanthology.org/2021.acl-long.434>

**AUTHORS:** Seanie Lee, Minki Kang, Juho Lee, Sung Ju Hwang

**HIGHLIGHT:** To tackle this problem, we propose a simple yet effective DA method based on a stochastic noise generator, which learns to perturb the word embedding of the input questions and context without changing their semantics.

435, **TITLE:** Maria: A Visual Experience Powered Conversational Agent

<https://aclanthology.org/2021.acl-long.435>

**AUTHORS:** Zujie Liang, Huang Hu, Can Xu, Chongyang Tao, Xiubo Geng, Yining Chen, Fan Liang, Daxin Jiang

**HIGHLIGHT:** In this paper, we take a step further to study image-grounded conversation under a fully open-ended setting where no paired dialog and image are assumed available.

436, **TITLE:** A Human-machine Collaborative Framework for Evaluating Malevolence in Dialogues

<https://aclanthology.org/2021.acl-long.436>

**AUTHORS:** Yangjun Zhang, Pengjie Ren, Maarten de Rijke

**HIGHLIGHT:** We introduce a human-machine collaborative framework, HMCEval, that can guarantee reliability of the evaluation outcomes with reduced human effort.

437, **TITLE:** Generating Relevant and Coherent Dialogue Responses using Self-Separated Conditional Variational AutoEncoders

<https://aclanthology.org/2021.acl-long.437>

**AUTHORS:** Bin Sun, Shaoxiong Feng, Yiwei Li, Jiamou Liu, Kan Li

**HIGHLIGHT:** To resolve this problem, we propose Self-separated Conditional Variational AutoEncoder (abbreviated as SepaCVAE) that introduces group information to regularize the latent variables, which enhances CVAE by improving the responses' relevance and coherence while maintaining their diversity and informativeness.

438, **TITLE:** Learning to Ask Conversational Questions by Optimizing Levenshtein Distance

<https://aclanthology.org/2021.acl-long.438>

**AUTHORS:** Zhongkun Liu, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Maarten de Rijke, Ming Zhou

**HIGHLIGHT:** In this work, we introduce a Reinforcement Iterative Sequence Editing (RISE) framework that optimizes the minimum Levenshtein distance through explicit editing actions.

439, **TITLE:** DVD: A Diagnostic Dataset for Multi-step Reasoning in Video Grounded Dialogue

<https://aclanthology.org/2021.acl-long.439>

**AUTHORS:** Hung Le, Chinnadhurai Sankar, Seungwhan Moon, Ahmad Beirami, Alborz Geramifard, Satwik Kottur

**HIGHLIGHT:** To address these limitations, in this paper, we present DVD, a Diagnostic Dataset for Video-grounded Dialogue.

- 440, TITLE: MMGCN: Multimodal Fusion via Deep Graph Convolution Network for Emotion Recognition in Conversation  
<https://aclanthology.org/2021.acl-long.440>  
AUTHORS: Jingwen Hu, Yuchen Liu, Jinming Zhao, Qin Jin  
HIGHLIGHT: In order to explore a more effective way of utilizing both multimodal and long-distance contextual information, we propose a new model based on multimodal fused graph convolutional network, MMGCN, in this work.
- 441, TITLE: DynaEval: Unifying Turn and Dialogue Level Evaluation  
<https://aclanthology.org/2021.acl-long.441>  
AUTHORS: Chen Zhang, Yiming Chen, Luis Fernando D'Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, Haizhou Li  
HIGHLIGHT: To this end, we propose DynaEval, a unified automatic evaluation framework which is not only capable of performing turn-level evaluation, but also holistically considers the quality of the entire dialogue.
- 442, TITLE: CoSQA: 20,000+ Web Queries for Code Search and Question Answering  
<https://aclanthology.org/2021.acl-long.442>  
AUTHORS: Junjie Huang, Duyu Tang, Linjun Shou, Ming Gong, Ke Xu, Daxin Jiang, Ming Zhou, Nan Duan  
HIGHLIGHT: To remedy this, we introduce CoSQA dataset. It includes 20,604 labels for pairs of natural language queries and codes, each annotated by at least 3 human annotators.
- 443, TITLE: Rewriter-Evaluator Architecture for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.443>  
AUTHORS: Yangming Li, Kaisheng Yao  
HIGHLIGHT: To address this issue, we introduce a novel architecture of Rewriter-Evaluator.
- 444, TITLE: Modeling Bilingual Conversational Characteristics for Neural Chat Translation  
<https://aclanthology.org/2021.acl-long.444>  
AUTHORS: Yunlong Liang, Fandong Meng, Yufeng Chen, Jinan Xu, Jie Zhou  
HIGHLIGHT: In this paper, we aim to promote the translation quality of conversational text by modeling the above properties.
- 445, TITLE: Importance-based Neuron Allocation for Multilingual Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.445>  
AUTHORS: Wanying Xie, Yang Feng, Shuhao Gu, Dong Yu  
HIGHLIGHT: To solve these problems, we propose to divide the model neurons into general and language-specific parts based on their importance across languages.
- 446, TITLE: Transfer Learning for Sequence Generation: from Single-source to Multi-source  
<https://aclanthology.org/2021.acl-long.446>  
AUTHORS: Xuancheng Huang, Jingfang Xu, Maosong Sun, Yang Liu  
HIGHLIGHT: Therefore, we propose a two-stage finetuning method to alleviate the pretrain-finetune discrepancy and introduce a novel MSG model with a fine encoder to learn better representations in MSG tasks.
- 447, TITLE: A Closer Look at Few-Shot Crosslingual Transfer: The Choice of Shots Matters  
<https://aclanthology.org/2021.acl-long.447>  
AUTHORS: Mengjie Zhao, Yi Zhu, Ehsan Shareghi, Ivan Vulic, Roi Reichart, Anna Korhonen, Hinrich Schütze  
HIGHLIGHT: In this work, we highlight a fundamental risk posed by this shortcoming, illustrating that the model exhibits a high degree of sensitivity to the selection of few shots.
- 448, TITLE: Coreference Reasoning in Machine Reading Comprehension  
<https://aclanthology.org/2021.acl-long.448>  
AUTHORS: Mingzhu Wu, Nafise Sadat Moosavi, Dan Roth, Iryna Gurevych  
HIGHLIGHT: We propose a methodology for creating MRC datasets that better reflect the challenges of coreference reasoning and use it to create a sample evaluation set.
- 449, TITLE: Adapting Unsupervised Syntactic Parsing Methodology for Discourse Dependency Parsing  
<https://aclanthology.org/2021.acl-long.449>  
AUTHORS: Liwen Zhang, Ge Wang, Wenjuan Han, Kewei Tu  
HIGHLIGHT: In this paper, we propose a simple yet effective method to adapt unsupervised syntactic dependency parsing methodology for unsupervised discourse dependency parsing.
- 450, TITLE: A Conditional Splitting Framework for Efficient Constituency Parsing

<https://aclanthology.org/2021.acl-long.450>

AUTHORS: Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty, Xiaoli Li  
HIGHLIGHT: We introduce a generic seq2seq parsing framework that casts constituency parsing problems (syntactic and discourse parsing) into a series of conditional splitting decisions.

451, TITLE: A Unified Generative Framework for Various NER Subtasks

<https://aclanthology.org/2021.acl-long.451>

AUTHORS: Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, Xipeng Qiu  
HIGHLIGHT: To that end, we propose to formulate the NER subtasks as an entity span sequence generation task, which can be solved by a unified sequence-to-sequence (Seq2Seq) framework.

452, TITLE: An In-depth Study on Internal Structure of Chinese Words

<https://aclanthology.org/2021.acl-long.452>

AUTHORS: Chen Gong, Saihao Huang, Houquan Zhou, Zhenghua Li, Min Zhang, Zhefeng Wang, Baoxing Huai, Nicholas Jing Yuan  
HIGHLIGHT: This work proposes to model the deep internal structures of Chinese words as dependency trees with 11 labels for distinguishing syntactic relationships.

453, TITLE: MulDA: A Multilingual Data Augmentation Framework for Low-Resource Cross-Lingual NER

<https://aclanthology.org/2021.acl-long.453>

AUTHORS: Linlin Liu, Bosheng Ding, Lidong Bing, Shafiq Joty, Luo Si, Chunyan Miao  
HIGHLIGHT: This paper addresses zero-shot transfer for cross-lingual NER, especially when the amount of source-language training data is also limited.

454, TITLE: Lexicon Enhanced Chinese Sequence Labeling Using BERT Adapter

<https://aclanthology.org/2021.acl-long.454>

AUTHORS: Wei Liu, Xiyan Fu, Yue Zhang, Wenming Xiao  
HIGHLIGHT: In this paper, we propose Lexicon Enhanced BERT (LEBERT) for Chinese sequence labeling, which integrates external lexicon knowledge into BERT layers directly by a Lexicon Adapter layer.

455, TITLE: Math Word Problem Solving with Explicit Numerical Values

<https://aclanthology.org/2021.acl-long.455>

AUTHORS: Qinzhuo Wu, Qi Zhang, Zhongyu Wei, Xuanjing Huang  
HIGHLIGHT: In this paper, we propose a novel approach called NumS2T, which enhances math word problem solving performance by explicitly incorporating numerical values into a sequence-to-tree network.

456, TITLE: Neural-Symbolic Solver for Math Word Problems with Auxiliary Tasks

<https://aclanthology.org/2021.acl-long.456>

AUTHORS: Jinghui Qin, Xiaodan Liang, Yining Hong, Jianheng Tang, Liang Lin  
HIGHLIGHT: Herein, we propose Neural-Symbolic Solver (NS-Solver) to explicitly and seamlessly incorporate different levels of symbolic constraints by auxiliary tasks.

457, TITLE: SMedBERT: A Knowledge-Enhanced Pre-trained Language Model with Structured Semantics for Medical Text Mining

<https://aclanthology.org/2021.acl-long.457>

AUTHORS: Taolin Zhang, Zerui Cai, Chengyu Wang, Minghui Qiu, Bite Yang, Xiaofeng He  
HIGHLIGHT: In this work, we introduce SMedBERT, a medical PLM trained on large-scale medical corpora, incorporating deep structured semantic knowledge from neighbours of linked-entity.

458, TITLE: What is Your Article Based On? Inferring Fine-grained Provenance

<https://aclanthology.org/2021.acl-long.458>

AUTHORS: Yi Zhang, Zachary Ives, Dan Roth  
HIGHLIGHT: In this paper, we introduce new techniques to model and reason about the provenance of multiple interacting claims, including how to capture fine-grained information about the context.

459, TITLE: Cross-modal Memory Networks for Radiology Report Generation

<https://aclanthology.org/2021.acl-long.459>

AUTHORS: Zhihong Chen, Yaling Shen, Yan Song, Xiang Wan  
HIGHLIGHT: In this paper, we propose a cross-modal memory networks (CMN) to enhance the encoder-decoder framework for radiology report generation, where a shared memory is designed to record the alignment between images and texts so as to facilitate the interaction and generation across modalities.

- 460, TITLE: Controversy and Conformity: from Generalized to Personalized Aggressiveness Detection  
<https://aclanthology.org/2021.acl-long.460>  
AUTHORS: Kamil Kanclerz, Alicja Figas, Marcin Gruza, Tomasz Kajdanowicz, Jan Kocon, Daria Puchalska, Przemyslaw Kazienko  
HIGHLIGHT: Therefore, we propose novel personalized approaches that respect individual beliefs expressed by either user conformity-based measures or various embeddings of their previous text annotations.
- 461, TITLE: Multi-perspective Coherent Reasoning for Helpfulness Prediction of Multimodal Reviews  
<https://aclanthology.org/2021.acl-long.461>  
AUTHORS: Junhao Liu, Zhen Hai, Min Yang, Lidong Bing  
HIGHLIGHT: This paper proposes a new task Multimodal Review Helpfulness Prediction (MRHP) aiming to analyze the review helpfulness from text and visual modalities.
- 462, TITLE: Instantaneous Grammatical Error Correction with Shallow Aggressive Decoding  
<https://aclanthology.org/2021.acl-long.462>  
AUTHORS: Xin Sun, Tao Ge, Furu Wei, Houfeng Wang  
HIGHLIGHT: In this paper, we propose Shallow Aggressive Decoding (SAD) to improve the online inference efficiency of the Transformer for instantaneous Grammatical Error Correction (GEC).
- 463, TITLE: Automatic ICD Coding via Interactive Shared Representation Networks with Self-distillation Mechanism  
<https://aclanthology.org/2021.acl-long.463>  
AUTHORS: Tong Zhou, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Kun Niu, Weifeng Chong, Shengping Liu  
HIGHLIGHT: To address the above issues, we propose an Interactive Shared Representation Network with Self-Distillation Mechanism.
- 464, TITLE: PHMOSpell: Phonological and Morphological Knowledge Guided Chinese Spelling Check  
<https://aclanthology.org/2021.acl-long.464>  
AUTHORS: Li Huang, Junjie Li, Weiwei Jiang, Zhiyu Zhang, Minchuan Chen, Shaojun Wang, Jing Xiao  
HIGHLIGHT: To address the above issues, we propose a novel end-to-end trainable model called PHMOSpell, which promotes the performance of CSC with multi-modal information.
- 465, TITLE: Guiding the Growth: Difficulty-Controllable Question Generation through Step-by-Step Rewriting  
<https://aclanthology.org/2021.acl-long.465>  
AUTHORS: Yi Cheng, Siyao Li, Bang Liu, Ruihui Zhao, Sujian Li, Chenghua Lin, Yefeng Zheng  
HIGHLIGHT: To this end, we propose a novel framework that progressively increases question difficulty through step-by-step rewriting under the guidance of an extracted reasoning chain.
- 466, TITLE: Improving Encoder by Auxiliary Supervision Tasks for Table-to-Text Generation  
<https://aclanthology.org/2021.acl-long.466>  
AUTHORS: Liang Li, Can Ma, Yinliang Yue, Dayong Hu  
HIGHLIGHT: Consequently, we propose to utilize two auxiliary tasks, Number Ranking (NR) and Importance Ranking (IR), to supervise the encoder to capture the different relations.
- 467, TITLE: POS-Constrained Parallel Decoding for Non-autoregressive Generation  
<https://aclanthology.org/2021.acl-long.467>  
AUTHORS: Kexin Yang, Wenqiang Lei, Dayiheng Liu, Weizhen Qi, Jiancheng Lv  
HIGHLIGHT: To provide a feasible solution to the multimodality problem of NAG, we propose incorporating linguistic structure (Part-of-Speech sequence in particular) into NAG inference instead of relying on teacher AG.
- 468, TITLE: Bridging Subword Gaps in Pretrain-Finetune Paradigm for Natural Language Generation  
<https://aclanthology.org/2021.acl-long.468>  
AUTHORS: Xin Liu, Baosong Yang, Dayiheng Liu, Haibo Zhang, Weihua Luo, Min Zhang, Haiying Zhang, Jinsong Su  
HIGHLIGHT: Towards approaching this problem, we extend the vanilla pretrain-finetune pipeline with an extra embedding transfer step.
- 469, TITLE: TGEA: An Error-Annotated Dataset and Benchmark Tasks for Text Generation from Pretrained Language Models  
<https://aclanthology.org/2021.acl-long.469>  
AUTHORS: Jie He, Bo Peng, Yi Liao, Qun Liu, Deyi Xiong



**HIGHLIGHT:** In order to deeply understand the capability of pretrained language models in text generation and conduct a diagnostic evaluation, we propose TGEA, an error-annotated dataset with multiple benchmark tasks for text generation from pretrained language models (PLMs).

470, **TITLE:** Long-Span Summarization via Local Attention and Content Selection

<https://aclanthology.org/2021.acl-long.470>

**AUTHORS:** Potsawee Manakul, Mark Gales

**HIGHLIGHT:** In this work, we exploit large pre-trained transformer-based models and address long-span dependencies in abstractive summarization using two methods: local self-attention; and explicit content selection.

471, **TITLE:** RepSum: Unsupervised Dialogue Summarization based on Replacement Strategy

<https://aclanthology.org/2021.acl-long.471>

**AUTHORS:** Xiyan Fu, Yating Zhang, Tianyi Wang, Xiaozhong Liu, Changlong Sun, Zhenglu Yang

**HIGHLIGHT:** In this study, we propose a novel unsupervised strategy to address this challenge, which roots from the hypothetical foundation that a superior summary approximates a replacement of the original dialogue, and they are roughly equivalent for auxiliary (self-supervised) tasks, e.g., dialogue generation.

472, **TITLE:** BASS: Boosting Abstractive Summarization with Unified Semantic Graph

<https://aclanthology.org/2021.acl-long.472>

**AUTHORS:** Wenhao Wu, Wei Li, Xinyan Xiao, Jiachen Liu, Ziqiang Cao, Sujian Li, Hua Wu, Haifeng Wang

**HIGHLIGHT:** In this paper, we present BASS, a novel framework for Boosting Abstractive Summarization based on a unified Semantic graph, which aggregates co-referent phrases distributing across a long range of context and conveys rich relations between phrases.

473, **TITLE:** Capturing Relations between Scientific Papers: An Abstractive Model for Related Work Section Generation

<https://aclanthology.org/2021.acl-long.473>

**AUTHORS:** Xiuying Chen, Hind Alamro, Mingzhe Li, Shen Gao, Xiangliang Zhang, Dongyan Zhao, Rui Yan

**HIGHLIGHT:** Hence, in this paper, we propose a Relation-aware Related work Generator (RRG), which generates an abstractive related work from the given multiple scientific papers in the same research area.

474, **TITLE:** Focus Attention: Promoting Faithfulness and Diversity in Summarization

<https://aclanthology.org/2021.acl-long.474>

**AUTHORS:** Rahul Aralikatte, Shashi Narayan, Joshua Maynez, Sascha Rothe, Ryan McDonald

**HIGHLIGHT:** With the motivation to narrow this gap, we introduce Focus Attention Mechanism, a simple yet effective method to encourage decoders to proactively generate tokens that are similar or topical to the input document.

475, **TITLE:** Generating Query Focused Summaries from Query-Free Resources

<https://aclanthology.org/2021.acl-long.475>

**AUTHORS:** Yumo Xu, Mirella Lapata

**HIGHLIGHT:** We introduce MaRGE, a Masked ROUGE Regression framework for evidence estimation and ranking which relies on a unified representation for summaries and queries, so that summaries in generic data can be converted into proxy queries for learning a query model.

476, **TITLE:** Robustifying Multi-hop QA through Pseudo-Evidentiality Training

<https://aclanthology.org/2021.acl-long.476>

**AUTHORS:** Kyungjae Lee, Seung-won Hwang, Sang-eun Han, Dohyeon Lee

**HIGHLIGHT:** In contrast, we propose a new approach to learn evidentiality, deciding whether the answer prediction is supported by correct evidences, without such annotations.

477, **TITLE:** xMoCo: Cross Momentum Contrastive Learning for Open-Domain Question Answering

<https://aclanthology.org/2021.acl-long.477>

**AUTHORS:** Nan Yang, Furu Wei, Binxing Jiao, Daxing Jiang, Linjun Yang

**HIGHLIGHT:** In this paper, we propose a new contrastive learning method called Cross Momentum Contrastive learning (xMoCo), for learning a dual-encoder model for question-passage matching.

478, **TITLE:** Learn to Resolve Conversational Dependency: A Consistency Training Framework for Conversational Question Answering

<https://aclanthology.org/2021.acl-long.478>

**AUTHORS:** Gangwoo Kim, Hyunjae Kim, Jungsoo Park, Jaewoo Kang

**HIGHLIGHT:** In this paper, we propose a novel framework, ExCorD (Explicit guidance on how to resolve Conversational Dependency) to enhance the abilities of QA models in comprehending conversational context.

479, TITLE: PhotoChat: A Human-Human Dialogue Dataset With Photo Sharing Behavior For Joint Image-Text Modeling  
<https://aclanthology.org/2021.acl-long.479>  
AUTHORS: Xiaoxue Zang, Lijuan Liu, Maria Wang, Yang Song, Hao Zhang, Jindong Chen  
HIGHLIGHT: We present a new human-human dialogue dataset - PhotoChat, the first dataset that casts light on the photo sharing behavior in online messaging.

480, TITLE: Good for Misconceived Reasons: An Empirical Revisiting on the Need for Visual Context in Multimodal Machine Translation  
<https://aclanthology.org/2021.acl-long.480>  
AUTHORS: Zhiyong Wu, Lingpeng Kong, Wei Bi, Xiang Li, Ben Kao  
HIGHLIGHT: Upon further investigation, we discover that the improvements achieved by the multimodal models over text-only counterparts are in fact results of the regularization effect.

481, TITLE: Attend What You Need: Motion-Appearance Synergistic Networks for Video Question Answering  
<https://aclanthology.org/2021.acl-long.481>  
AUTHORS: Ahjeong Seo, Gi-Cheon Kang, Joonhan Park, Byoung-Tak Zhang  
HIGHLIGHT: We propose Motion-Appearance Synergistic Networks (MASN), which embed two cross-modal features grounded on motion and appearance information and selectively utilize them depending on the question's intentions.

482, TITLE: BERTifying the Hidden Markov Model for Multi-Source Weakly Supervised Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.482>  
AUTHORS: Yinghao Li, Pranav Shetty, Lucas Liu, Chao Zhang, Le Song  
HIGHLIGHT: To address this challenge, we propose a conditional hidden Markov model (CHMM), which can effectively infer true labels from multi-source noisy labels in an unsupervised way.

483, TITLE: CIL: Contrastive Instance Learning Framework for Distantly Supervised Relation Extraction  
<https://aclanthology.org/2021.acl-long.483>  
AUTHORS: Tao Chen, Haizhou Shi, Siliang Tang, Zhigang Chen, Fei Wu, Yueting Zhuang  
HIGHLIGHT: In this paper, we go beyond typical MIL framework and propose a novel contrastive instance learning (CIL) framework.

484, TITLE: SENT: Sentence-level Distant Relation Extraction via Negative Training  
<https://aclanthology.org/2021.acl-long.484>  
AUTHORS: Ruotian Ma, Tao Gui, Linyang Li, Qi Zhang, Xuanjing Huang, Yaqian Zhou  
HIGHLIGHT: Based on NT, we propose a sentence-level framework, SENT, for distant relation extraction.

485, TITLE: An End-to-End Progressive Multi-Task Learning Framework for Medical Named Entity Recognition and Normalization  
<https://aclanthology.org/2021.acl-long.485>  
AUTHORS: Baohang Zhou, Xiangrui Cai, Ying Zhang, Xiaojie Yuan  
HIGHLIGHT: To avoid the disadvantages of existing models and exploit the generalized representation across the two tasks, we design an end-to-end progressive multi-task learning model for jointly modeling medical NER and NEN in an effective way.

486, TITLE: PRGC: Potential Relation and Global Correspondence Based Joint Relational Triple Extraction  
<https://aclanthology.org/2021.acl-long.486>  
AUTHORS: Hengyi Zheng, Rui Wen, Xi Chen, Yifan Yang, Yunyan Zhang, Ziheng Zhang, Ningyu Zhang, Bin Qin, Xu Ming, Yefeng Zheng  
HIGHLIGHT: In this paper, we decompose this task into three subtasks, Relation Judgement, Entity Extraction and Subject-object Alignment from a novel perspective and then propose a joint relational triple extraction framework based on Potential Relation and Global Correspondence (PRGC).

487, TITLE: Learning from Miscellaneous Other-Class Words for Few-shot Named Entity Recognition  
<https://aclanthology.org/2021.acl-long.487>  
AUTHORS: Meihan Tong, Shuai Wang, Bin Xu, Yixin Cao, Minghui Liu, Lei Hou, Juanzi Li  
HIGHLIGHT: To address the issue, we propose a novel model, Mining Undefined Classes from Other-class (MUCO), that can automatically induce different unde- fined classes from the other class to improve few-shot NER.

488, TITLE: Joint Biomedical Entity and Relation Extraction with Knowledge-Enhanced Collective Inference  
<https://aclanthology.org/2021.acl-long.488>

- AUTHORS: Tuan Lai, Heng Ji, ChengXiang Zhai, Quan Hung Tran  
HIGHLIGHT: Inspired by how humans look up relevant information to comprehend a scientific text, we present a novel framework that utilizes external knowledge for joint entity and relation extraction named KECI (Knowledge-Enhanced Collective Inference).
- 489, TITLE: Fine-grained Information Extraction from Biomedical Literature based on Knowledge-enriched Abstract Meaning Representation  
<https://aclanthology.org/2021.acl-long.489>  
AUTHORS: Zixuan Zhang, Nikolaus Parulian, Heng Ji, Ahmed Elsayed, Skatje Myers, Martha Palmer  
HIGHLIGHT: In this paper, we propose a novel biomedical Information Extraction (IE) model to tackle these two challenges and extract scientific entities and events from English research papers.
- 490, TITLE: Unleash GPT-2 Power for Event Detection  
<https://aclanthology.org/2021.acl-long.490>  
AUTHORS: Amir Pouran Ben Veysch, Viet Lai, Franck Dernoncourt, Thien Huu Nguyen  
HIGHLIGHT: To overcome this issue, we propose to exploit the powerful pre-trained language model GPT-2 to generate training samples for ED.
- 491, TITLE: CLEVE: Contrastive Pre-training for Event Extraction  
<https://aclanthology.org/2021.acl-long.491>  
AUTHORS: Ziqi Wang, Xiaozhi Wang, Xu Han, Yankai Lin, Lei Hou, Zhiyuan Liu, Peng Li, Juanzi Li, Jie Zhou  
HIGHLIGHT: To this end, we propose CLEVE, a contrastive pre-training framework for EE to better learn event knowledge from large unsupervised data and their semantic structures (e.g. AMR) obtained with automatic parsers.
- 492, TITLE: Document-level Event Extraction via Parallel Prediction Networks  
<https://aclanthology.org/2021.acl-long.492>  
AUTHORS: Hang Yang, Dianbo Sui, Yubo Chen, Kang Liu, Jun Zhao, Taifeng Wang  
HIGHLIGHT: In this paper, we propose an end-to-end model, which can extract structured events from a document in a parallel manner.
- 493, TITLE: StructuralLM: Structural Pre-training for Form Understanding  
<https://aclanthology.org/2021.acl-long.493>  
AUTHORS: Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, Luo Si  
HIGHLIGHT: In this paper, we propose a new pre-training approach, StructuralLM, to jointly leverage cell and layout information from scanned documents.
- 494, TITLE: Dual Graph Convolutional Networks for Aspect-based Sentiment Analysis  
<https://aclanthology.org/2021.acl-long.494>  
AUTHORS: Ruihan Li, Hao Chen, Fangxiang Feng, Zhanyu Ma, Xiaojie Wang, Eduard Hovy  
HIGHLIGHT: To overcome these challenges, in this paper, we propose a dual graph convolutional networks (DualGCN) model that considers the complementarity of syntax structures and semantic correlations simultaneously.
- 495, TITLE: Multi-Label Few-Shot Learning for Aspect Category Detection  
<https://aclanthology.org/2021.acl-long.495>  
AUTHORS: Mengting Hu, Shiwan Zhao, Honglei Guo, Chao Xue, Hang Gao, Tiegang Gao, Renhong Cheng, Zhong Su  
HIGHLIGHT: In this paper, we formulate ACD in the few-shot learning scenario.
- 496, TITLE: Argument Pair Extraction via Attention-guided Multi-Layer Multi-Cross Encoding  
<https://aclanthology.org/2021.acl-long.496>  
AUTHORS: Liying Cheng, Tianyu Wu, Lidong Bing, Luo Si  
HIGHLIGHT: This paper proposes a novel attention-guided multi-layer multi-cross encoding scheme to address the challenges.
- 497, TITLE: A Neural Transition-based Model for Argumentation Mining  
<https://aclanthology.org/2021.acl-long.497>  
AUTHORS: Jianzhu Bao, Chuang Fan, Jipeng Wu, Yixue Dang, Jiachen Du, Ruifeng Xu  
HIGHLIGHT: Towards these issues, we propose a neural transition-based model for argumentation mining, which incrementally builds an argumentation graph by generating a sequence of actions, avoiding inefficient enumeration operations.
- 498, TITLE: Keep It Simple: Unsupervised Simplification of Multi-Paragraph Text

- <https://aclanthology.org/2021.acl-long.498>  
AUTHORS: Philippe Laban, Tobias Schnabel, Paul Bennett, Marti A. Hearst  
HIGHLIGHT: This work presents Keep it Simple (KiS), a new approach to unsupervised text simplification which learns to balance a reward across three properties: fluency, salience and simplicity.
- 499, TITLE: Long Text Generation by Modeling Sentence-Level and Discourse-Level Coherence  
<https://aclanthology.org/2021.acl-long.499>  
AUTHORS: Jian Guan, Xiaoxi Mao, Changjie Fan, Zitao Liu, Wenbiao Ding, Minlie Huang  
HIGHLIGHT: In this paper, we propose a long text generation model, which can represent the prefix sentences at sentence level and discourse level in the decoding process.
- 500, TITLE: OpenMEVA: A Benchmark for Evaluating Open-ended Story Generation Metrics  
<https://aclanthology.org/2021.acl-long.500>  
AUTHORS: Jian Guan, Zhixin Zhang, Zhuoer Feng, Zitao Liu, Wenbiao Ding, Xiaoxi Mao, Changjie Fan, Minlie Huang  
HIGHLIGHT: OpenMEVA provides a comprehensive test suite to assess the capabilities of metrics, including (a) the correlation with human judgments, (b) the generalization to different model outputs and datasets, (c) the ability to judge story coherence, and (d) the robustness to perturbations
- 501, TITLE: DYPLOC: Dynamic Planning of Content Using Mixed Language Models for Text Generation  
<https://aclanthology.org/2021.acl-long.501>  
AUTHORS: Xinyu Hua, Ashwin Sreevatsa, Lu Wang  
HIGHLIGHT: To this end, we propose DYPLOC, a generation framework that conducts dynamic planning of content while generating the output based on a novel design of mixed language models.
- 502, TITLE: Controllable Open-ended Question Generation with A New Question Type Ontology  
<https://aclanthology.org/2021.acl-long.502>  
AUTHORS: Shuyang Cao, Lu Wang  
HIGHLIGHT: We investigate the less-explored task of generating open-ended questions that are typically answered by multiple sentences.
- 503, TITLE: BERTGen: Multi-task Generation through BERT  
<https://aclanthology.org/2021.acl-long.503>  
AUTHORS: Faidon Mitzalis, Ozan Caglayan, Pranava Madhyastha, Lucia Specia  
HIGHLIGHT: We present BERTGen, a novel, generative, decoder-only model which extends BERT by fusing multimodal and multilingual pre-trained models VL-BERT and M-BERT, respectively.
- 504, TITLE: Selective Knowledge Distillation for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.504>  
AUTHORS: Fusheng Wang, Jianhao Yan, Fandong Meng, Jie Zhou  
HIGHLIGHT: In this paper, we design a novel protocol that can effectively analyze the different impacts of samples by comparing various samples' partitions.
- 505, TITLE: Measuring and Increasing Context Usage in Context-Aware Machine Translation  
<https://aclanthology.org/2021.acl-long.505>  
AUTHORS: Patrick Fernandes, Kayo Yin, Graham Neubig, Andr  F. T. Martins  
HIGHLIGHT: In this paper, we introduce a new metric, conditional cross-mutual information, to quantify usage of context by these models.
- 506, TITLE: Beyond Offline Mapping: Learning Cross-lingual Word Embeddings through Context Anchoring  
<https://aclanthology.org/2021.acl-long.506>  
AUTHORS: Aitor Ormazabal, Mikel Artetxe, Aitor Soroa, Gorka Labaka, Eneko Agirre  
HIGHLIGHT: In this paper, we propose an alternative approach that does not have this limitation, while requiring a weak seed dictionary (e.g., a list of identical words) as the only form of supervision.
- 507, TITLE: CCMatrix: Mining Billions of High-Quality Parallel Sentences on the Web  
<https://aclanthology.org/2021.acl-long.507>  
AUTHORS: Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, Angela Fan  
HIGHLIGHT: We show that margin-based bitext mining in a multilingual sentence space can be successfully scaled to operate on monolingual corpora of billions of sentences.

- 508, TITLE: Length-Adaptive Transformer: Train Once with Length Drop, Use Anytime with Search  
<https://aclanthology.org/2021.acl-long.508>  
AUTHORS: Gyuwan Kim, Kyunghyun Cho  
HIGHLIGHT: In this paper, we extend PoWER-BERT (Goyal et al., 2020) and propose Length-Adaptive Transformer that can be used for various inference scenarios after one-shot training.
- 509, TITLE: GhostBERT: Generate More Features with Cheap Operations for BERT  
<https://aclanthology.org/2021.acl-long.509>  
AUTHORS: Zhiqi Huang, Lu Hou, Lifeng Shang, Xin Jiang, Xiao Chen, Qun Liu  
HIGHLIGHT: In this paper, we propose GhostBERT, which generates more features with very cheap operations from the remaining features.
- 510, TITLE: Super Tickets in Pre-Trained Language Models: From Model Compression to Improving Generalization  
<https://aclanthology.org/2021.acl-long.510>  
AUTHORS: Chen Liang, Simiao Zuo, Minshuo Chen, Haoming Jiang, Xiaodong Liu, Pengcheng He, Tuo Zhao, Weizhu Chen  
HIGHLIGHT: In this paper, we study such a collection of tickets, which is referred to as "winning tickets", in extremely over-parametrized models, e.g., pre-trained language models.
- 511, TITLE: A Novel Estimator of Mutual Information for Learning to Disentangle Textual Representations  
<https://aclanthology.org/2021.acl-long.511>  
AUTHORS: Pierre Colombo, Pablo Piantanida, Chloé Clavel  
HIGHLIGHT: This paper introduces a novel variational upper bound to the mutual information between an attribute and the latent code of an encoder.
- 512, TITLE: Determinantal Beam Search  
<https://aclanthology.org/2021.acl-long.512>  
AUTHORS: Clara Meister, Martina Forster, Ryan Cotterell  
HIGHLIGHT: To address this issue, we propose a reformulation of beam search, which we call determinantal beam search.
- 513, TITLE: Multi-hop Graph Convolutional Network with High-order Chebyshev Approximation for Text Reasoning  
<https://aclanthology.org/2021.acl-long.513>  
AUTHORS: Shuoran Jiang, Qingcai Chen, Xin Liu, Baotian Hu, Lisai Zhang  
HIGHLIGHT: In this study, we define the spectral graph convolutional network with the high-order dynamic Chebyshev approximation (HDGCN), which augments the multi-hop graph reasoning by fusing messages aggregated from direct and long-term dependencies into one convolutional layer.
- 514, TITLE: Accelerating Text Communication via Abbreviated Sentence Input  
<https://aclanthology.org/2021.acl-long.514>  
AUTHORS: Jiban Adhikary, Jamie Berger, Keith Vertanen  
HIGHLIGHT: We designed a recognizer optimized for expanding noisy abbreviated input where users often omit spaces and mid-word vowels.
- 515, TITLE: Regression Bugs Are In Your Model! Measuring, Reducing and Analyzing Regressions In NLP Model Updates  
<https://aclanthology.org/2021.acl-long.515>  
AUTHORS: Yuqing Xie, Yi-An Lai, Yuanjun Xiong, Yi Zhang, Stefano Soatto  
HIGHLIGHT: This work focuses on quantifying, reducing and analyzing regression errors in the NLP model updates.
- 516, TITLE: Detecting Propaganda Techniques in Memes  
<https://aclanthology.org/2021.acl-long.516>  
AUTHORS: Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, Giovanni Da San Martino  
HIGHLIGHT: With this in mind, here we propose a new multi-label multimodal task: detecting the type of propaganda techniques used in memes.
- 517, TITLE: On the Efficacy of Adversarial Data Collection for Question Answering: Results from a Large-Scale Randomized Study  
<https://aclanthology.org/2021.acl-long.517>  
AUTHORS: Divyansh Kaushik, Douwe Kiela, Zachary C. Lipton, Wen-tau Yih  
HIGHLIGHT: In this paper, we conduct a large-scale controlled study focused on question answering, assigning workers at random to compose questions either (i) adversarially (with a model in the loop); or (ii) in the standard fashion (without a model).

- 518, TITLE: Learning Dense Representations of Phrases at Scale  
<https://aclanthology.org/2021.acl-long.518>  
AUTHORS: Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, Danqi Chen  
HIGHLIGHT: We present an effective method to learn phrase representations from the supervision of reading comprehension tasks, coupled with novel negative sampling methods.
- 519, TITLE: End-to-End Training of Neural Retrievers for Open-Domain Question Answering  
<https://aclanthology.org/2021.acl-long.519>  
AUTHORS: Devendra Sachan, Mostofa Patwary, Mohammad Shoeybi, Neel Kant, Wei Ping, William L. Hamilton, Bryan Catanzaro  
HIGHLIGHT: In this work, we systematically study retriever pre-training.
- 520, TITLE: Question Answering Over Temporal Knowledge Graphs  
<https://aclanthology.org/2021.acl-long.520>  
AUTHORS: Apoorv Saxena, Soumen Chakrabarti, Partha Talukdar  
HIGHLIGHT: We address this challenge by presenting CRONQUESTIONS, the largest known Temporal KGQA dataset, clearly stratified into buckets of structural complexity.
- 521, TITLE: Language Model Augmented Relevance Score  
<https://aclanthology.org/2021.acl-long.521>  
AUTHORS: Ruibo Liu, Jason Wei, Soroush Vosoughi  
HIGHLIGHT: In this paper, we propose Language Model Augmented Relevance Score (MARS), a new context-aware metric for NLG evaluation.
- 522, TITLE: DExperts: Decoding-Time Controlled Text Generation with Experts and Anti-Experts  
<https://aclanthology.org/2021.acl-long.522>  
AUTHORS: Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, Yejin Choi  
HIGHLIGHT: We propose DExperts: Decoding-time Experts, a decoding-time method for controlled text generation that combines a pretrained language model with "expert" LMs and/or "anti-expert" LMs in a product of experts.
- 523, TITLE: Polyjuice: Generating Counterfactuals for Explaining, Evaluating, and Improving Models  
<https://aclanthology.org/2021.acl-long.523>  
AUTHORS: Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, Daniel Weld  
HIGHLIGHT: We present Polyjuice, a general-purpose counterfactual generator that allows for control over perturbation types and locations, trained by finetuning GPT-2 on multiple datasets of paired sentences.
- 524, TITLE: Metaphor Generation with Conceptual Mappings  
<https://aclanthology.org/2021.acl-long.524>  
AUTHORS: Kevin Stowe, Tuhin Chakrabarty, Nanyun Peng, Smaranda Muresan, Iryna Gurevych  
HIGHLIGHT: In this paper, we aim to generate a metaphoric sentence given a literal expression by replacing relevant verbs.
- 525, TITLE: Learning Latent Structures for Cross Action Phrase Relations in Wet Lab Protocols  
<https://aclanthology.org/2021.acl-long.525>  
AUTHORS: Chaitanya Kulkarni, Jany Chan, Eric Fosler-Lussier, Raghu Machiraju  
HIGHLIGHT: Here, we propose methods to automatically generate a MSTG for a given protocol by extracting all action relationships across multiple sentences.
- 526, TITLE: Multimodal Multi-Speaker Merger & Acquisition Financial Modeling: A New Task, Dataset, and Neural Baselines  
<https://aclanthology.org/2021.acl-long.526>  
AUTHORS: Ramit Sawhney, Mihir Goyal, Prakhar Goel, Puneet Mathur, Rajiv Ratn Shah  
HIGHLIGHT: We introduce M3ANet, a baseline architecture that takes advantage of the multimodal multi-speaker input to forecast the financial risk associated with the M&A calls.
- 527, TITLE: Mid-Air Hand Gestures for Post-Editing of Machine Translation  
<https://aclanthology.org/2021.acl-long.527>  
AUTHORS: Rashad Albo Jamara, Nico Herbig, Antonio Kr?ger, Josef van Genabith

**HIGHLIGHT:** Here, we present the first study that investigates the usefulness of mid-air hand gestures in combination with the keyboard (GK) for text editing in PE of MT. Guided by a gesture elicitation study with 14 freelance translators, we develop a prototype supporting mid-air hand gestures for cursor placement, text selection, deletion, and reordering.

528, **TITLE:** Inter-GPS: Interpretable Geometry Problem Solving with Formal Language and Symbolic Reasoning  
<https://aclanthology.org/2021.acl-long.528>

**AUTHORS:** Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, Song-Chun Zhu  
**HIGHLIGHT:** Thus, we construct a new large-scale benchmark, Geometry3K, consisting of 3,002 geometry problems with dense annotation in formal language. We further propose a novel geometry solving approach with formal language and symbolic reasoning, called Interpretable Geometry Problem Solver (Inter-GPS).

529, **TITLE:** Joint Verification and Reranking for Open Fact Checking Over Tables  
<https://aclanthology.org/2021.acl-long.529>

**AUTHORS:** Michael Sejr Schlichtkrull, Vladimir Karpukhin, Barlas Oguz, Mike Lewis, Wen-tau Yih, Sebastian Riedel  
**HIGHLIGHT:** In this paper, we investigate verification over structured data in the open-domain setting, introducing a joint reranking-and-verification model which fuses evidence documents in the verification component.

530, **TITLE:** Evaluation of Thematic Coherence in Microblogs  
<https://aclanthology.org/2021.acl-long.530>

**AUTHORS:** Iman Munire Bilal, Bo Wang, Maria Liakata, Rob Procter, Adam Tsakalidis  
**HIGHLIGHT:** Here we create a corpus of microblog clusters from three different domains and time windows and define the task of evaluating thematic coherence. We provide annotation guidelines and human annotations of thematic coherence by journalist experts.

531, **TITLE:** Neural semi-Markov CRF for Monolingual Word Alignment  
<https://aclanthology.org/2021.acl-long.531>

**AUTHORS:** Wuwei Lan, Chao Jiang, Wei Xu  
**HIGHLIGHT:** In this paper, we present a novel neural semi-Markov CRF alignment model, which unifies word and phrase alignments through variable-length spans.

532, **TITLE:** Privacy at Scale: Introducing the PrivaSeer Corpus of Web Privacy Policies  
<https://aclanthology.org/2021.acl-long.532>

**AUTHORS:** Mukund Srinath, Shomir Wilson, C Lee Giles  
**HIGHLIGHT:** We describe a corpus creation pipeline with stages that include a web crawler, language detection, document classification, duplicate and near-duplicate removal, and content extraction.

533, **TITLE:** The statistical advantage of automatic NLG metrics at the system level  
<https://aclanthology.org/2021.acl-long.533>

**AUTHORS:** Johnny Wei, Robin Jia  
**HIGHLIGHT:** This paper qualifies the notion that automatic metrics are not as good as humans in estimating system-level quality.

534, **TITLE:** Are Missing Links Predictable? An Inferential Benchmark for Knowledge Graph Completion  
<https://aclanthology.org/2021.acl-long.534>

**AUTHORS:** Yixin Cao, Xiang Ji, Xin Lv, Juanzi Li, Yonggang Wen, Hanwang Zhang  
**HIGHLIGHT:** We present InferWiki, a Knowledge Graph Completion (KGC) dataset that improves upon existing benchmarks in inferential ability, assumptions, and patterns.

535, **TITLE:** ConvoSumm: Conversation Summarization Benchmark and Improved Abstractive Summarization with Argument Mining  
<https://aclanthology.org/2021.acl-long.535>

**AUTHORS:** Alexander Fabbri, Faiaz Rahman, Imad Rizvi, Borui Wang, Haoran Li, Yashar Mehdad, Dragomir Radev  
**HIGHLIGHT:** To address this gap, we design annotation protocols motivated by an issues-viewpoints-assertions framework to crowdsource four new datasets on diverse online conversation forms of news comments, discussion forums, community question answering forums, and email threads.

536, **TITLE:** Improving Factual Consistency of Abstractive Summarization via Question Answering  
<https://aclanthology.org/2021.acl-long.536>

**AUTHORS:** Feng Nan, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Kathleen McKeown, Ramesh Nallapati, Dejjiao Zhang, Zhiguo Wang, Andrew O. Arnold, Bing Xiang  
**HIGHLIGHT:** In this paper we present an approach to address factual consistency in summarization.

- 537, TITLE: EmailSum: Abstractive Email Thread Summarization  
<https://aclanthology.org/2021.acl-long.537>  
AUTHORS: Shiyue Zhang, Asli Celikyilmaz, Jianfeng Gao, Mohit Bansal  
HIGHLIGHT: To spur research in thread summarization, we have developed an abstractive Email Thread Summarization (EmailSum) dataset, which contains human-annotated short (<30 words) and long (<100 words) summaries of 2,549 email threads (each containing 3 to 10 emails) over a wide variety of topics.
- 538, TITLE: Cross-Lingual Abstractive Summarization with Limited Parallel Resources  
<https://aclanthology.org/2021.acl-long.538>  
AUTHORS: Yu Bai, Yang Gao, Heyan Huang  
HIGHLIGHT: To bridge these connections, we propose a novel Multi-Task framework for Cross-Lingual Abstractive Summarization (MCLAS) in a low-resource setting.
- 539, TITLE: Dissecting Generation Modes for Abstractive Summarization Models via Ablation and Attribution  
<https://aclanthology.org/2021.acl-long.539>  
AUTHORS: Jiacheng Xu, Greg Durrett  
HIGHLIGHT: We propose a two-step method to interpret summarization model decisions.
- 540, TITLE: Learning Prototypical Functions for Physical Artifacts  
<https://aclanthology.org/2021.acl-long.540>  
AUTHORS: Tianyu Jiang, Ellen Riloff  
HIGHLIGHT: In this paper, we introduce a new NLP task of learning the prototypical uses for human-made physical objects.
- 541, TITLE: Verb Knowledge Injection for Multilingual Event Processing  
<https://aclanthology.org/2021.acl-long.541>  
AUTHORS: Olga Majewska, Ivan Vulic, Goran Glava?, Edoardo Maria Ponti, Anna Korhonen  
HIGHLIGHT: In this paper, we target a specific facet of linguistic knowledge, the interplay between verb meaning and argument structure.
- 542, TITLE: Dynamic Contextualized Word Embeddings  
<https://aclanthology.org/2021.acl-long.542>  
AUTHORS: Valentin Hofmann, Janet Pierrehumbert, Hinrich Sch?tze  
HIGHLIGHT: Building on prior work on contextualized and dynamic word embeddings, we introduce dynamic contextualized word embeddings that represent words as a function of both linguistic and extralinguistic context.
- 543, TITLE: Lexical Semantic Change Discovery  
<https://aclanthology.org/2021.acl-long.543>  
AUTHORS: Sinan Kurtuyigit, Maïke Park, Dominik Schlechtweg, Jonas Kuhn, Sabine Schulte im Walde  
HIGHLIGHT: In this paper, we propose a shift of focus from change detection to change discovery, i.e., discovering novel word senses over time from the full corpus vocabulary.
- 544, TITLE: The R-U-A-Robot Dataset: Helping Avoid Chatbot Deception by Detecting User Questions About Human or Non-Human Identity  
<https://aclanthology.org/2021.acl-long.544>  
AUTHORS: David Gros, Yu Li, Zhou Yu  
HIGHLIGHT: We collect over 2,500 phrasings related to the intent of "Are you a robot?". This is paired with over 2,500 adversarially selected utterances where only confirming the system is non-human would be insufficient or disfluent.
- 545, TITLE: Using Meta-Knowledge Mined from Identifiers to Improve Intent Recognition in Conversational Systems  
<https://aclanthology.org/2021.acl-long.545>  
AUTHORS: Claudio Pinhanez, Paulo Cavalin, Victor Henrique Alves Ribeiro, Ana Appel, Heloisa Candello, Julio Nogima, Mauro Pichiliani, Melina Guerra, Maira de Bayser, Gabriel Malfatti, Henrique Ferreira  
HIGHLIGHT: In this paper we explore the improvement of intent recognition in conversational systems by the use of meta-knowledge embedded in intent identifiers.
- 546, TITLE: Space Efficient Context Encoding for Non-Task-Oriented Dialogue Generation with Graph Attention Transformer  
<https://aclanthology.org/2021.acl-long.546>  
AUTHORS: Fabian Galetzka, Jewgeni Rose, David Schlangen, Jens Lehmann



**HIGHLIGHT:** In this work, we propose a more concise encoding for background context structured in the form of knowledge graphs, by expressing the graph connections through restrictions on the attention weights.

547, **TITLE:** DialogueCRN: Contextual Reasoning Networks for Emotion Recognition in Conversations  
<https://aclanthology.org/2021.acl-long.547>

**AUTHORS:** Dou Hu, Lingwei Wei, Xiaoyong Huai

**HIGHLIGHT:** In this work, we propose novel Contextual Reasoning Networks (DialogueCRN) to fully understand the conversational context from a cognitive perspective.

548, **TITLE:** Cross-replication Reliability - An Empirical Approach to Interpreting Inter-rater Reliability  
<https://aclanthology.org/2021.acl-long.548>

**AUTHORS:** Ka Wong, Praveen Paritosh, Lora Aroyo

**HIGHLIGHT:** We present a new alternative to interpreting IRR that is more empirical and contextualized.

549, **TITLE:** TIMEDIAL: Temporal Commonsense Reasoning in Dialog

<https://aclanthology.org/2021.acl-long.549>

**AUTHORS:** Lianhui Qin, Aditya Gupta, Shyam Upadhyay, Luheng He, Yejin Choi, Manaal Faruqui

**HIGHLIGHT:** In this paper, we present the first study to investigate pre-trained LMs for their temporal reasoning capabilities in dialogs by introducing a new task and a crowd-sourced English challenge set, TimeDial.

550, **TITLE:** RAW-C: Relatedness of Ambiguous Words in Context (A New Lexical Resource for English)

<https://aclanthology.org/2021.acl-long.550>

**AUTHORS:** Sean Trott, Benjamin Bergen

**HIGHLIGHT:** We introduce RAW-C, a dataset of graded, human relatedness judgments for 112 ambiguous words in context (with 672 sentence pairs total), as well as human estimates of sense dominance.

551, **TITLE:** ARBERT & MARBERT: Deep Bidirectional Transformers for Arabic

<https://aclanthology.org/2021.acl-long.551>

**AUTHORS:** Muhammad Abdul-Mageed, AbdelRahim Elmadany, El Moatez Billah Nagoudi

**HIGHLIGHT:** Although multilingual LMs were also introduced to serve many languages, these have limitations such as being costly at inference time and the size and diversity of non-English data involved in their pre-training. We remedy these issues for a collection of diverse Arabic varieties by introducing two powerful deep bidirectional transformer-based models, ARBERT and MARBERT.

552, **TITLE:** Improving Paraphrase Detection with the Adversarial Paraphrasing Task

<https://aclanthology.org/2021.acl-long.552>

**AUTHORS:** Animesh Nigohjkar, John Licato

**HIGHLIGHT:** We apply the adversarial paradigm to this question, and introduce a new adversarial method of dataset creation for paraphrase identification: the Adversarial Paraphrasing Task (APT), which asks participants to generate semantically equivalent (in the sense of mutually implicative) but lexically and syntactically disparate paraphrases.

553, **TITLE:** ADEPT: An Adjective-Dependent Plausibility Task

<https://aclanthology.org/2021.acl-long.553>

**AUTHORS:** Ali Emami, Ian Porada, Alexandra Olteanu, Kaheer Suleman, Adam Trischler, Jackie Chi Kit Cheung

**HIGHLIGHT:** We introduce ADEPT - a large-scale semantic plausibility task consisting of over 16 thousand sentences that are paired with slightly modified versions obtained by adding an adjective to a noun.

554, **TITLE:** ReadOnce Transformers: Reusable Representations of Text for Transformers

<https://aclanthology.org/2021.acl-long.554>

**AUTHORS:** Shih-Ting Lin, Ashish Sabharwal, Tushar Khot

**HIGHLIGHT:** We present ReadOnce Transformers, an approach to convert a transformer-based model into one that can build an information-capturing, task-independent, and compressed representation of text.

555, **TITLE:** Conditional Generation of Temporally-ordered Event Sequences

<https://aclanthology.org/2021.acl-long.555>

**AUTHORS:** Shih-Ting Lin, Nathanael Chambers, Greg Durrett

**HIGHLIGHT:** We propose a single model that addresses both temporal ordering, sorting given events into the order they occurred, and event infilling, predicting new events which fit into an existing temporally-ordered sequence.

556, **TITLE:** Hate Speech Detection Based on Sentiment Knowledge Sharing

- <https://aclanthology.org/2021.acl-long.556>  
AUTHORS: Xianbing Zhou, Yang Yong, Xiaochao Fan, Ge Ren, Yunfeng Song, Yufeng Diao, Liang Yang, Hongfei Lin  
HIGHLIGHT: In this paper, we propose a hate speech detection framework based on sentiment knowledge sharing.
- 557, TITLE: Transition-based Bubble Parsing: Improvements on Coordination Structure Prediction  
<https://aclanthology.org/2021.acl-long.557>  
AUTHORS: Tianze Shi, Lillian Lee  
HIGHLIGHT: In this paper, we introduce a transition system and neural models for parsing these bubble-enhanced structures.
- 558, TITLE: SpanNER: Named Entity Re-/Recognition as Span Prediction  
<https://aclanthology.org/2021.acl-long.558>  
AUTHORS: Jinlan Fu, Xuanjing Huang, Pengfei Liu  
HIGHLIGHT: In this paper, we first investigate the strengths and weaknesses when the span prediction model is used for named entity recognition compared with the sequence labeling framework and how to further improve it, which motivates us to make complementary advantages of systems based on different paradigms. We then reveal that span prediction, simultaneously, can serve as a system combiner to re-recognize named entities from different systems' outputs.
- 559, TITLE: StructFormer: Joint Unsupervised Induction of Dependency and Constituency Structure from Masked Language Modeling  
<https://aclanthology.org/2021.acl-long.559>  
AUTHORS: Yikang Shen, Yi Tay, Che Zheng, Dara Bahri, Donald Metzler, Aaron Courville  
HIGHLIGHT: To achieve this, we propose a new parsing framework that can jointly generate a constituency tree and dependency graph.
- 560, TITLE: Language Embeddings for Typology and Cross-lingual Transfer Learning  
<https://aclanthology.org/2021.acl-long.560>  
AUTHORS: Dian Yu, Taiqi He, Kenji Sagae  
HIGHLIGHT: We explore whether language representations that capture relationships among languages can be learned and subsequently leveraged in cross-lingual tasks without the use of parallel data.
- 561, TITLE: Can Sequence-to-Sequence Models Crack Substitution Ciphers?  
<https://aclanthology.org/2021.acl-long.561>  
AUTHORS: Nada Aldarrab, Jonathan May  
HIGHLIGHT: We propose an end-to-end multilingual model for solving simple substitution ciphers.
- 562, TITLE: Beyond Noise: Mitigating the Impact of Fine-grained Semantic Divergences on Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.562>  
AUTHORS: Eleftheria Briakou, Marine Carpuat  
HIGHLIGHT: Based on these findings, we introduce a divergent-aware NMT framework that uses factors to help NMT recover from the degradation caused by naturally occurring divergences, improving both translation quality and model calibration on EN-FR tasks.
- 563, TITLE: Discriminative Reranking for Neural Machine Translation  
<https://aclanthology.org/2021.acl-long.563>  
AUTHORS: Ann Lee, Michael Auli, Marc?Aurelio Ranzato  
HIGHLIGHT: Since such a discriminator contains hundreds of millions of parameters, we improve its generalization using pre-training and data augmentation techniques.
- 564, TITLE: Mind Your Outliers! Investigating the Negative Impact of Outliers on Active Learning for Visual Question Answering  
<https://aclanthology.org/2021.acl-long.564>  
AUTHORS: Siddharth Karamcheti, Ranjay Krishna, Li Fei-Fei, Christopher Manning  
HIGHLIGHT: To understand this discrepancy, we profile 8 active learning methods on a per-example basis, and identify the problem as collective outliers - groups of examples that active learning methods prefer to acquire but models fail to learn (e.g., questions that ask about text in images or require external knowledge).
- 565, TITLE: All That's 'Human' Is Not Gold: Evaluating Human Evaluation of Generated Text  
<https://aclanthology.org/2021.acl-long.565>  
AUTHORS: Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, Noah A. Smith

**HIGHLIGHT:** We explore three approaches for quickly training evaluators to better identify GPT3-authored text (detailed instructions, annotated examples, and paired examples) and find that while evaluators' accuracy improved up to 55%, it did not significantly improve across the three domains.

566, **TITLE:** Scientific Credibility of Machine Translation Research: A Meta-Evaluation of 769 Papers  
<https://aclanthology.org/2021.acl-long.566>

**AUTHORS:** Benjamin Marie, Atsushi Fujita, Raphael Rubino

**HIGHLIGHT:** This paper presents the first large-scale meta-evaluation of machine translation (MT).

567, **TITLE:** Neural Machine Translation with Monolingual Translation Memory

<https://aclanthology.org/2021.acl-long.567>

**AUTHORS:** Deng Cai, Yan Wang, Huayang Li, Wai Lam, Lema Liu

**HIGHLIGHT:** In contrast to existing work that uses bilingual corpus as TM and employs source-side similarity search for memory retrieval, we propose a new framework that uses monolingual memory and performs learnable memory retrieval in a cross-lingual manner.

568, **TITLE:** Intrinsic Dimensionality Explains the Effectiveness of Language Model Fine-Tuning

<https://aclanthology.org/2021.acl-long.568>

**AUTHORS:** Armen Aghajanyan, Sonal Gupta, Luke Zettlemoyer

**HIGHLIGHT:** In this paper, we argue that analyzing fine-tuning through the lens of intrinsic dimension provides us with empirical and theoretical intuitions to explain this remarkable phenomenon.

569, **TITLE:** UnNatural Language Inference

<https://aclanthology.org/2021.acl-long.569>

**AUTHORS:** Koustuv Sinha, Prasanna Parthasarathi, Joelle Pineau, Adina Williams

**HIGHLIGHT:** To measure the severity of this issue, we propose a suite of metrics and investigate which properties of particular permutations lead models to be word order invariant.

570, **TITLE:** Including Signed Languages in Natural Language Processing

<https://aclanthology.org/2021.acl-long.570>

**AUTHORS:** Kayo Yin, Amit Moryossef, Julie Hochgesang, Yoav Goldberg, Malihe Alikhani

**HIGHLIGHT:** This position paper calls on the NLP community to include signed languages as a research area with high social and scientific impact.

571, **TITLE:** Vocabulary Learning via Optimal Transport for Neural Machine Translation

<https://aclanthology.org/2021.acl-long.571>

**AUTHORS:** Jingjing Xu, Hao Zhou, Chun Gan, Zaixiang Zheng, Lei Li

**HIGHLIGHT:** This paper aims to figure out what is a good vocabulary and whether we can find the optimal vocabulary without trial training.

572, **TITLE:** Catchphrase: Automatic Detection of Cultural References

<https://aclanthology.org/2021.acl-short.1>

**AUTHORS:** Nir Sweed, Dafna Shahaf

**HIGHLIGHT:** In this paper, we study snowclones originating from pop-culture quotes; our goal is to automatically detect cultural references in text.

573, **TITLE:** On Training Instance Selection for Few-Shot Neural Text Generation

<https://aclanthology.org/2021.acl-short.2>

**AUTHORS:** Ernie Chang, Xiaoyu Shen, Hui-Syuan Yeh, Vera Demberg

**HIGHLIGHT:** In this work, we present a study on training instance selection in few-shot neural text generation.

574, **TITLE:** Coreference Resolution without Span Representations

<https://aclanthology.org/2021.acl-short.3>

**AUTHORS:** Yuval Kirstain, Ori Ram, Omer Levy

**HIGHLIGHT:** We introduce a lightweight end-to-end coreference model that removes the dependency on span representations, handcrafted features, and heuristics.

575, **TITLE:** Enhancing Entity Boundary Detection for Better Chinese Named Entity Recognition

<https://aclanthology.org/2021.acl-short.4>

**AUTHORS:** Chun Chen, Fang Kong

HIGHLIGHT: In this paper, we propose a boundary enhanced approach for better Chinese NER.

576, TITLE: Difficulty-Aware Machine Translation Evaluation

<https://aclanthology.org/2021.acl-short.5>

AUTHORS: Runzhe Zhan, Xuebo Liu, Derek F. Wong, Lidia S. Chao

HIGHLIGHT: In this paper, we propose a novel difficulty-aware MT evaluation metric, expanding the evaluation dimension by taking translation difficulty into consideration.

577, TITLE: Uncertainty and Surprisal Jointly Deliver the Punchline: Exploiting Incongruity-Based Features for Humor Recognition

<https://aclanthology.org/2021.acl-short.6>

AUTHORS: Yubo Xie, Junze Li, Pearl Pu

HIGHLIGHT: Inspired by the incongruity theory of humor, we model the set-up as the part developing semantic uncertainty, and the punchline disrupting audience expectations.

578, TITLE: Counterfactuals to Control Latent Disentangled Text Representations for Style Transfer

<https://aclanthology.org/2021.acl-short.7>

AUTHORS: Sharmila Reddy Nangi, Niyati Chhaya, Sopan Khosla, Nikhil Kaushik, Harshit Nyati

HIGHLIGHT: In this paper, we propose a counterfactual-based method to modify the latent representation, by posing a 'what-if' scenario.

579, TITLE: Attention Flows are Shapley Value Explanations

<https://aclanthology.org/2021.acl-short.8>

AUTHORS: Kawin Ethayarajh, Dan Jurafsky

HIGHLIGHT: Attention flow is a post-processed variant of attention weights obtained by running the max-flow algorithm on the attention graph. Perhaps surprisingly, we prove that attention flows are indeed Shapley Values, at least at the layerwise level.

580, TITLE: Video Paragraph Captioning as a Text Summarization Task

<https://aclanthology.org/2021.acl-short.9>

AUTHORS: Hui Liu, Xiaojun Wan

HIGHLIGHT: In this work, we propose a novel framework by taking this task as a text summarization task.

581, TITLE: Are VQA Systems RAD? Measuring Robustness to Augmented Data with Focused Interventions

<https://aclanthology.org/2021.acl-short.10>

AUTHORS: Daniel Rosenberg, Itai Gat, Amir Feder, Roi Reichart

HIGHLIGHT: Using these augmentations, we propose a new robustness measure, Robustness to Augmented Data (RAD), which measures the consistency of model predictions between original and augmented examples.

582, TITLE: How Helpful is Inverse Reinforcement Learning for Table-to-Text Generation?

<https://aclanthology.org/2021.acl-short.11>

AUTHORS: Sayan Ghosh, Zheng Qi, Snigdha Chaturvedi, Shashank Srivastava

HIGHLIGHT: In this work, we instead pose the Table-to-Text task as Inverse Reinforcement Learning (IRL) problem.

583, TITLE: Automatic Fake News Detection: Are Models Learning to Reason?

<https://aclanthology.org/2021.acl-short.12>

AUTHORS: Casper Hansen, Christian Hansen, Lucas Chaves Lima

HIGHLIGHT: In this paper, we investigate this assumption of reasoning, by exploring the relationship and importance of both claim and evidence.

584, TITLE: Saying No is An Art: Contextualized Fallback Responses for Unanswerable Dialogue Queries

<https://aclanthology.org/2021.acl-short.13>

AUTHORS: Ashish Shrivastava, Kaustubh Dhole, Abhinav Bhatt, Sharvani Raghunath

HIGHLIGHT: While, dialogue systems today rely on static and unnatural responses like "I don't know the answer to that question" or "I'm not sure about that", we design a neural approach which generates responses which are contextually aware with the user query as well as say no to the user

585, TITLE: N-Best ASR Transformer: Enhancing SLU Performance using Multiple ASR Hypotheses

<https://aclanthology.org/2021.acl-short.14>

AUTHORS: Karthik Ganesan, Pakhi Bamdev, Jaivarsan B, Amresh Venugopal, Abhinav Tushar

**HIGHLIGHT:** We hypothesize that transformer models will learn better with a simpler utterance representation using the concatenation of the N-best ASR alternatives, where each alternative is separated by a special delimiter [SEP].

586, **TITLE:** Gender bias amplification during Speed-Quality optimization in Neural Machine Translation

<https://aclanthology.org/2021.acl-short.15>

**AUTHORS:** Adithya Renduchintala, Denise Diaz, Kenneth Heafield, Xian Li, Mona Diab

**HIGHLIGHT:** We investigate architectures and techniques commonly used to speed up decoding in Transformer-based models, such as greedy search, quantization, average attention networks (AANs) and shallow decoder models and show their effect on gendered noun translation.

587, **TITLE:** Machine Translation into Low-resource Language Varieties

<https://aclanthology.org/2021.acl-short.16>

**AUTHORS:** Sachin Kumar, Antonios Anastasopoulos, Shuly Wintner, Yulia Tsvetkov

**HIGHLIGHT:** We propose a general framework to rapidly adapt MT systems to generate language varieties that are close to, but different from, the standard target language, using no parallel (source-variety) data.

588, **TITLE:** Is Sparse Attention more Interpretable?

<https://aclanthology.org/2021.acl-short.17>

**AUTHORS:** Clara Meister, Stefan Lazov, Isabelle Augenstein, Ryan Cotterell

**HIGHLIGHT:** We build on the recent work exploring the interpretability of attention; we design a set of experiments to help us understand how sparsity affects our ability to use attention as an explainability tool.

589, **TITLE:** The Case for Translation-Invariant Self-Attention in Transformer-Based Language Models

<https://aclanthology.org/2021.acl-short.18>

**AUTHORS:** Ulme Wennberg, Gustav Eje Henter

**HIGHLIGHT:** In this paper, we analyze the position embeddings of existing language models, finding strong evidence of translation invariance, both for the embeddings themselves and for their effect on self-attention.

590, **TITLE:** Relative Importance in Sentence Processing

<https://aclanthology.org/2021.acl-short.19>

**AUTHORS:** Nora Hollenstein, Lisa Beinborn

**HIGHLIGHT:** In this work, we compare patterns of relative importance in English language processing by humans and models and analyze the underlying linguistic patterns.

591, **TITLE:** Doing Good or Doing Right? Exploring the Weakness of Commonsense Causal Reasoning Models

<https://aclanthology.org/2021.acl-short.20>

**AUTHORS:** Mingyue Han, Yinglin Wang

**HIGHLIGHT:** In this paper, we investigate the problem of semantic similarity bias and reveal the vulnerability of current COPA models by certain attacks.

592, **TITLE:** AND does not mean OR: Using Formal Languages to Study Language Models? Representations

<https://aclanthology.org/2021.acl-short.21>

**AUTHORS:** Aaron Traylor, Roman Feiman, Ellie Pavlick

**HIGHLIGHT:** The goal of this study is to offer insights into a narrower but critical subquestion: Under what conditions should we expect that meaning and form covary sufficiently, such that a language model with access only to form might nonetheless succeed in emulating meaning?

593, **TITLE:** Enforcing Consistency in Weakly Supervised Semantic Parsing

<https://aclanthology.org/2021.acl-short.22>

**AUTHORS:** Nitish Gupta, Sameer Singh, Matt Gardner

**HIGHLIGHT:** In this work we explore the use of consistency between the output programs for related inputs to reduce the impact of spurious programs.

594, **TITLE:** An Improved Model for Voicing Silent Speech

<https://aclanthology.org/2021.acl-short.23>

**AUTHORS:** David Gaddy, Dan Klein

**HIGHLIGHT:** In this paper, we present an improved model for voicing silent speech, where audio is synthesized from facial electromyography (EMG) signals.

595, **TITLE:** What's in the Box? An Analysis of Undesirable Content in the Common Crawl Corpus

- <https://aclanthology.org/2021.acl-short.24>  
AUTHORS: Alexandra Luccioni, Joseph Viviano  
HIGHLIGHT: We find that it contains a significant amount of undesirable content, including hate speech and sexually explicit content, even after filtering procedures.
- 596, TITLE: Continual Quality Estimation with Online Bayesian Meta-Learning  
<https://aclanthology.org/2021.acl-short.25>  
AUTHORS: Abiola Obamuyide, Marina Fomicheva, Lucia Specia  
HIGHLIGHT: To address this challenge, we propose an online Bayesian meta-learning framework for the continuous training of QE models that is able to adapt them to the needs of different users, while being robust to distributional shifts in training and test data.
- 597, TITLE: A Span-based Dynamic Local Attention Model for Sequential Sentence Classification  
<https://aclanthology.org/2021.acl-short.26>  
AUTHORS: Xichen Shang, Qianli Ma, Zhenxi Lin, Jiangyue Yan, Zipeng Chen  
HIGHLIGHT: In this paper, we proposed a span-based dynamic local attention model that could explicitly capture the structural information by the proposed supervised dynamic local attention.
- 598, TITLE: How effective is BERT without word ordering? Implications for language understanding and data privacy  
<https://aclanthology.org/2021.acl-short.27>  
AUTHORS: Jack Hessel, Alexandra Schofield  
HIGHLIGHT: We show that the token representations and self-attention activations within BERT are surprisingly resilient to shuffling the order of input tokens, and that for several GLUE language understanding tasks, shuffling only minimally degrades performance, e.g., by 4% for QNLI.
- 599, TITLE: WikiSum: Coherent Summarization Dataset for Efficient Human-Evaluation  
<https://aclanthology.org/2021.acl-short.28>  
AUTHORS: Nachshon Cohen, Oren Kalinsky, Yftah Ziser, Alessandro Moschitti  
HIGHLIGHT: To overcome these limitations, we present a dataset based on article summaries appearing on the WikiHow website, composed of how-to articles and coherent-paragraph summaries written in plain language.
- 600, TITLE: UMIC: An Unreferenced Metric for Image Captioning via Contrastive Learning  
<https://aclanthology.org/2021.acl-short.29>  
AUTHORS: Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Trung Bui, Kyomin Jung  
HIGHLIGHT: In this paper, we introduce a new metric UMIC, an Unreferenced Metric for Image Captioning which does not require reference captions to evaluate image captions.
- 601, TITLE: Anchor-based Bilingual Word Embeddings for Low-Resource Languages  
<https://aclanthology.org/2021.acl-short.30>  
AUTHORS: Tobias Eder, Viktor Hangya, Alexander Fraser  
HIGHLIGHT: This paper proposes a new approach for building BWEs in which the vector space of the high resource source language is used as a starting point for training an embedding space for the low resource target language.
- 602, TITLE: Multilingual Agreement for Multilingual Neural Machine Translation  
<https://aclanthology.org/2021.acl-short.31>  
AUTHORS: Jian Yang, Yuwei Yin, Shuming Ma, Haoyang Huang, Dongdong Zhang, Zhoujun Li, Furu Wei  
HIGHLIGHT: In this work, we propose a novel agreement-based method to encourage multilingual agreement among different translation directions, which minimizes the differences among them.
- 603, TITLE: Higher-order Derivatives of Weighted Finite-state Machines  
<https://aclanthology.org/2021.acl-short.32>  
AUTHORS: Ran Zmigrod, Tim Vieira, Ryan Cotterell  
HIGHLIGHT: This work examines the computation of higher-order derivatives with respect to the normalization constant for weighted finite-state machines.
- 604, TITLE: Reinforcement Learning for Abstractive Question Summarization with Question-aware Semantic Rewards  
<https://aclanthology.org/2021.acl-short.33>  
AUTHORS: Shweta Yadav, Deepak Gupta, Asma Ben Abacha, Dina Demner-Fushman  
HIGHLIGHT: In this paper, we introduce a reinforcement learning-based framework for abstractive question summarization.

- 605, TITLE: A Semantics-aware Transformer Model of Relation Linking for Knowledge Base Question Answering  
<https://aclanthology.org/2021.acl-short.34>  
AUTHORS: Tahira Naseem, Srinivas Ravishankar, Nandana Mihindukulasooriya, Ibrahim Abdelaziz, Young-Suk Lee, Pavan Kapanipathi, Salim Roukos, Alfio Gliozzo, Alexander Gray  
HIGHLIGHT: We propose a simple transformer-based neural model for relation linking that leverages the AMR semantic parse of a sentence.
- 606, TITLE: Neural Retrieval for Question Answering with Cross-Attention Supervised Data Augmentation  
<https://aclanthology.org/2021.acl-short.35>  
AUTHORS: Yinfei Yang, Ning Jin, Kuo Lin, Mandy Guo, Daniel Cer  
HIGHLIGHT: We present a supervised data mining method using an accurate early fusion model to improve the training of an efficient late fusion retrieval model.
- 607, TITLE: Enhancing Descriptive Image Captioning with Natural Language Inference  
<https://aclanthology.org/2021.acl-short.36>  
AUTHORS: Zhan Shi, Hui Liu, Xiaodan Zhu  
HIGHLIGHT: In this paper we propose a novel approach to encourage captioning models to produce more detailed captions using natural language inference, based on the motivation that, among different captions of an image, descriptive captions are more likely to entail less descriptive captions.
- 608, TITLE: MOLEMAN: Mention-Only Linking of Entities with a Mention Annotation Network  
<https://aclanthology.org/2021.acl-short.37>  
AUTHORS: Nicholas FitzGerald, Dan Bikel, Jan Botha, Daniel Gillick, Tom Kwiatkowski, Andrew McCallum  
HIGHLIGHT: We present an instance-based nearest neighbor approach to entity linking.
- 609, TITLE: eMLM: A New Pre-training Objective for Emotion Related Tasks  
<https://aclanthology.org/2021.acl-short.38>  
AUTHORS: Tiberiu Sosea, Cornelia Caragea  
HIGHLIGHT: In this paper, we present Emotion Masked Language Modelling, a variation of Masked Language Modelling aimed at improving the BERT language representation model for emotion detection and sentiment analysis tasks.
- 610, TITLE: On Positivity Bias in Negative Reviews  
<https://aclanthology.org/2021.acl-short.39>  
AUTHORS: Madhusudhan Aithal, Chenhao Tan  
HIGHLIGHT: Consistent with prior work, we show that English negative reviews tend to contain more positive words than negative words, using a variety of datasets.
- 611, TITLE: PRAL: A Tailored Pre-Training Model for Task-Oriented Dialog Generation  
<https://aclanthology.org/2021.acl-short.40>  
AUTHORS: Jing Gu, Qingyang Wu, Chongruo Wu, Weiyan Shi, Zhou Yu  
HIGHLIGHT: We propose a Pre-trained Role Alternating Language model (PRAL), explicitly designed for task-oriented conversational systems.
- 612, TITLE: ROPE: Reading Order Equivariant Positional Encoding for Graph-based Document Information Extraction  
<https://aclanthology.org/2021.acl-short.41>  
AUTHORS: Chen-Yu Lee, Chun-Liang Li, Chu Wang, Renshen Wang, Yasuhisa Fujii, Siyang Qin, Ashok Papat, Tomas Pfister  
HIGHLIGHT: We propose Reading Order Equivariant Positional Encoding (ROPE), a new positional encoding technique designed to apprehend the sequential presentation of words in documents.
- 613, TITLE: Zero-shot Event Extraction via Transfer Learning: Challenges and Insights  
<https://aclanthology.org/2021.acl-short.42>  
AUTHORS: Qing Lyu, Hongming Zhang, Elijor Sulem, Dan Roth  
HIGHLIGHT: In this work, we explore the possibility of zero-shot event extraction by formulating it as a set of Textual Entailment (TE) and/or Question Answering (QA) queries (e.g. "A city was attacked" entails "There is an attack"), exploiting pretrained TE/QA models for direct transfer.
- 614, TITLE: Using Adversarial Attacks to Reveal the Statistical Bias in Machine Reading Comprehension Models  
<https://aclanthology.org/2021.acl-short.43>  
AUTHORS: Jieyu Lin, Jiajie Zou, Nai Ding

- HIGHLIGHT:** Here, we demonstrate a simple yet effective method to attack MRC models and reveal the statistical biases in these models.
- 615, **TITLE:** Quantifying and Avoiding Unfair Qualification Labour in Crowdsourcing  
<https://aclanthology.org/2021.acl-short.44>  
**AUTHORS:** Jonathan K. Kummerfeld  
**HIGHLIGHT:** Through analysis of worker discussions and guidance for researchers, we estimate that workers spend approximately 2.25 months of full time effort on poorly paid tasks in order to get the qualifications needed for better paid tasks.
- 616, **TITLE:** Men Are Elected, Women Are Married: Events Gender Bias on Wikipedia  
<https://aclanthology.org/2021.acl-short.45>  
**AUTHORS:** Jiao Sun, Nanyun Peng  
**HIGHLIGHT:** In this paper, we present the first event-centric study of gender biases in a Wikipedia corpus.
- 617, **TITLE:** Modeling Task-Aware MIMO Cardinality for Efficient Multilingual Neural Machine Translation  
<https://aclanthology.org/2021.acl-short.46>  
**AUTHORS:** Hongfei Xu, Qiuhui Liu, Josef van Genabith, Deyi Xiong  
**HIGHLIGHT:** In this paper, we propose to efficiently increase the capacity for multilingual NMT by increasing the cardinality.
- 618, **TITLE:** Adaptive Nearest Neighbor Machine Translation  
<https://aclanthology.org/2021.acl-short.47>  
**AUTHORS:** Xin Zheng, Zhirui Zhang, Junliang Guo, Shujian Huang, Boxing Chen, Weihua Luo, Jiajun Chen  
**HIGHLIGHT:** In this paper, we propose Adaptive kNN-MT to dynamically determine the number of k for each target token.
- 619, **TITLE:** On Orthogonality Constraints for Transformers  
<https://aclanthology.org/2021.acl-short.48>  
**AUTHORS:** Aston Zhang, Alvin Chan, Yi Tay, Jie Fu, Shuohang Wang, Shuai Zhang, Huajie Shao, Shuochao Yao, Roy Ka-Wei Lee  
**HIGHLIGHT:** To fill this gap, this paper studies orthogonality constraints for transformers, showing the effectiveness with empirical evidence from ten machine translation tasks and two dialogue generation tasks.
- 620, **TITLE:** Measuring and Improving BERT 's Mathematical Abilities by Predicting the Order of Reasoning.  
<https://aclanthology.org/2021.acl-short.49>  
**AUTHORS:** Piotr Piekos, Mateusz Malinowski, Henryk Michalewski  
**HIGHLIGHT:** Towards this goal, we investigate if a commonly used language model, BERT, possesses such mathematical abilities and, if so, to what degree.
- 621, **TITLE:** Happy Dance, Slow Clap: Using Reaction GIFs to Predict Induced Affect on Twitter  
<https://aclanthology.org/2021.acl-short.50>  
**AUTHORS:** Boaz Shmueli, Soumya Ray, Lun-Wei Ku  
**HIGHLIGHT:** We present a new, automated method for collecting texts along with their induced reaction labels.
- 622, **TITLE:** Exploring Listwise Evidence Reasoning with T5 for Fact Verification  
<https://aclanthology.org/2021.acl-short.51>  
**AUTHORS:** Kelvin Jiang, Ronak Pradeep, Jimmy Lin  
**HIGHLIGHT:** This work explores a framework for fact verification that leverages pretrained sequence-to-sequence transformer models for sentence selection and label prediction, two key sub-tasks in fact verification.
- 623, **TITLE:** DefSent: Sentence Embeddings using Definition Sentences  
<https://aclanthology.org/2021.acl-short.52>  
**AUTHORS:** Hayato Tsukagoshi, Ryohei Sasano, Koichi Takeda  
**HIGHLIGHT:** In this paper, we propose DefSent, a sentence embedding method that uses definition sentences from a word dictionary, which performs comparably on unsupervised semantics textual similarity (STS) tasks and slightly better on SentEval tasks than conventional methods.
- 624, **TITLE:** Discrete Cosine Transform as Universal Sentence Encoder  
<https://aclanthology.org/2021.acl-short.53>  
**AUTHORS:** Nada Almarwani, Mona Diab  
**HIGHLIGHT:** To this end, we utilize DCT encoder to generate universal sentence representation for different languages such as German, French, Spanish and Russian.



- 625, TITLE: AligNarr: Aligning Narratives on Movies  
<https://aclanthology.org/2021.acl-short.54>  
AUTHORS: Paramita Mirza, Mostafa Abouhamra, Gerhard Weikum  
HIGHLIGHT: This paper addresses the alignment problem by devising a fully unsupervised approach based on a global optimization model.
- 626, TITLE: An Exploratory Analysis of Multilingual Word-Level Quality Estimation with Cross-Lingual Transformers  
<https://aclanthology.org/2021.acl-short.55>  
AUTHORS: Tharindu Ranasinghe, Constantin Orasan, Ruslan Mitkov  
HIGHLIGHT: To overcome these problems, we explore different approaches to multilingual, word-level QE.
- 627, TITLE: Exploration and Exploitation: Two Ways to Improve Chinese Spelling Correction Models  
<https://aclanthology.org/2021.acl-short.56>  
AUTHORS: Chong Li, Cenyuan Zhang, Xiaoqing Zheng, Xuanjing Huang  
HIGHLIGHT: We propose a method, which continually identifies the weak spots of a model to generate more valuable training instances, and apply a task-specific pre-training strategy to enhance the model.
- 628, TITLE: Training Adaptive Computation for Open-Domain Question Answering with Computational Constraints  
<https://aclanthology.org/2021.acl-short.57>  
AUTHORS: Yuxiang Wu, Pasquale Minervini, Pontus Stenertorp, Sebastian Riedel  
HIGHLIGHT: We propose Adaptive Passage Encoder, an AC method that can be applied to an existing ODQA model and can be trained efficiently on a single GPU.
- 629, TITLE: An Empirical Study on Adversarial Attack on NMT: Languages and Positions Matter  
<https://aclanthology.org/2021.acl-short.58>  
AUTHORS: Zhiyuan Zeng, Deyi Xiong  
HIGHLIGHT: In this paper, we empirically investigate adversarial attack on NMT from two aspects: languages (the source vs. the target language) and positions (front vs. rear).
- 630, TITLE: OntoGUM: Evaluating Contextualized SOTA Coreference Resolution on 12 More Genres  
<https://aclanthology.org/2021.acl-short.59>  
AUTHORS: Yilun Zhu, Sameer Pradhan, Amir Zeldes  
HIGHLIGHT: This paper provides a dataset and comprehensive evaluation showing that the latest neural LM based end-to-end systems degrade very substantially out of domain.
- 631, TITLE: In Factuality: Efficient Integration of Relevant Facts for Visual Question Answering  
<https://aclanthology.org/2021.acl-short.60>  
AUTHORS: Peter Vickers, Nikolaos Aletras, Emilio Monti, Loïc Barrault  
HIGHLIGHT: In this paper, we propose a new method to enhance the reasoning capabilities of a multi-modal pretrained model (Vision+Language BERT) by integrating facts extracted from an external knowledge base.
- 632, TITLE: Zero-shot Fact Verification by Claim Generation  
<https://aclanthology.org/2021.acl-short.61>  
AUTHORS: Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan, William Yang Wang  
HIGHLIGHT: We develop QACG, a framework for training a robust fact verification model by using automatically generated claims that can be supported, refuted, or unverifiable from evidence from Wikipedia.
- 633, TITLE: Thank you BART! Rewarding Pre-Trained Models Improves Formality Style Transfer  
<https://aclanthology.org/2021.acl-short.62>  
AUTHORS: Huiyuan Lai, Antonio Toral, Malvina Nissim  
HIGHLIGHT: We show that fine-tuning pre-trained language (GPT-2) and sequence-to-sequence (BART) models boosts content preservation, and that this is possible even with limited amounts of parallel data.
- 634, TITLE: Deep Context- and Relation-Aware Learning for Aspect-based Sentiment Analysis  
<https://aclanthology.org/2021.acl-short.63>  
AUTHORS: Shinhyeok Oh, Dongyub Lee, Taesun Whang, IlNam Park, Seo Gaeun, EungGyun Kim, Harksoo Kim  
HIGHLIGHT: In this paper, we propose Deep Contextualized Relation-Aware Network (DCRAN), which allows interactive relations among subtasks with deep contextual information based on two modules (i.e., Aspect and Opinion Propagation and Explicit Self-Supervised Strategies).

- 635, TITLE: Towards Generative Aspect-Based Sentiment Analysis  
<https://aclanthology.org/2021.acl-short.64>  
AUTHORS: Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, Wai Lam  
HIGHLIGHT: In this paper, we propose to tackle various ABSA tasks in a unified generative framework.
- 636, TITLE: Bilingual Mutual Information Based Adaptive Training for Neural Machine Translation  
<https://aclanthology.org/2021.acl-short.65>  
AUTHORS: Yangyifan Xu, Yijin Liu, Fandong Meng, Jiajun Zhang, Jinan Xu, Jie Zhou  
HIGHLIGHT: In this paper, we propose a novel bilingual mutual information (BMI) based adaptive objective, which measures the learning difficulty for each target token from the perspective of bilingualism, and assigns an adaptive weight accordingly to improve token-level adaptive training.
- 637, TITLE: Continual Learning for Task-oriented Dialogue System with Iterative Network Pruning, Expanding and Masking  
<https://aclanthology.org/2021.acl-short.66>  
AUTHORS: Binzong Geng, Fajie Yuan, Qiancheng Xu, Ying Shen, Ruifeng Xu, Min Yang  
HIGHLIGHT: This paper proposes an effective continual learning method for the task-oriented dialogue system with iterative network pruning, expanding, and masking (TPEM), which preserves performance on previously encountered tasks while accelerating learning progress on subsequent tasks.
- 638, TITLE: TIMERS: Document-level Temporal Relation Extraction  
<https://aclanthology.org/2021.acl-short.67>  
AUTHORS: Puneet Mathur, Rajiv Jain, Franck Dernoncourt, Vlad Morariu, Quan Hung Tran, Dinesh Manocha  
HIGHLIGHT: We present TIMERS - a TIME, Rhetorical and Syntactic-aware model for document-level temporal relation classification in the English language.
- 639, TITLE: Improving Arabic Diacritization with Regularized Decoding and Adversarial Training  
<https://aclanthology.org/2021.acl-short.68>  
AUTHORS: Han Qin, Guimin Chen, Yuanhe Tian, Yan Song  
HIGHLIGHT: In this paper, we propose to use regularized decoding and adversarial training to appropriately learn from such noisy knowledge for diacritization.
- 640, TITLE: When is Char Better Than Subword: A Systematic Study of Segmentation Algorithms for Neural Machine Translation  
<https://aclanthology.org/2021.acl-short.69>  
AUTHORS: Jiahuan Li, Yutong Shen, Shujian Huang, Xinyu Dai, Jiajun Chen  
HIGHLIGHT: In this paper, we present an in-depth comparison between character-based and subword-based NMT systems under three settings: translating to typologically diverse languages, training with low resource, and adapting to unseen domains.
- 641, TITLE: More than Text: Multi-modal Chinese Word Segmentation  
<https://aclanthology.org/2021.acl-short.70>  
AUTHORS: Dong Zhang, Zheng Hu, Shoushan Li, Hanqian Wu, Qiaoming Zhu, Guodong Zhou  
HIGHLIGHT: In this paper, we annotate a new dataset for CWS containing text and audio. Moreover, we propose a time-dependent multi-modal interactive model based on Transformer framework to integrate multi-modal information for word sequence labeling.
- 642, TITLE: A Mixture-of-Experts Model for Anonym-Synonym Discrimination  
<https://aclanthology.org/2021.acl-short.71>  
AUTHORS: Zhipeng Xie, Nan Zeng  
HIGHLIGHT: This paper proposes two underlying hypotheses and employs the mixture-of-experts framework as a solution.
- 643, TITLE: Learning Domain-Specialised Representations for Cross-Lingual Biomedical Entity Linking  
<https://aclanthology.org/2021.acl-short.72>  
AUTHORS: Fangyu Liu, Ivan Vulic, Anna Korhonen, Nigel Collier  
HIGHLIGHT: To this end, we propose and evaluate a series of cross-lingual transfer methods for the XL-BEL task, and demonstrate that general-domain bitext helps propagate the available English knowledge to languages with little to no in-domain data.
- 644, TITLE: A Cluster-based Approach for Improving Isotropy in Contextual Embedding Space  
<https://aclanthology.org/2021.acl-short.73>

- AUTHORS: Sara Rajaei, Mohammad Taher Pilehvar  
HIGHLIGHT: Based on this observation, we propose a local cluster-based method to address the degeneration issue in contextual embedding spaces.
- 645, TITLE: Unsupervised Enrichment of Persona-grounded Dialog with Background Stories  
<https://aclanthology.org/2021.acl-short.74>  
AUTHORS: Bodhisattwa Prasad Majumder, Taylor Berg-Kirkpatrick, Julian McAuley, Harsh Jhamtani  
HIGHLIGHT: In this work, we equip dialog models with 'background stories' related to a persona by leveraging fictional narratives from existing story datasets (e.g. ROCStories).
- 646, TITLE: Beyond Laurel/Yanny: An Autoencoder-Enabled Search for Polyperceivable Audio  
<https://aclanthology.org/2021.acl-short.75>  
AUTHORS: Kartik Chandra, Chuma Kabaghe, Gregory Valiant  
HIGHLIGHT: In this paper we apply ML techniques to study the prevalence of polyperceivability in spoken language.
- 647, TITLE: Don't Let Discourse Confine Your Model: Sequence Perturbations for Improved Event Language Models  
<https://aclanthology.org/2021.acl-short.76>  
AUTHORS: Mahnaz Koupaee, Greg Durrett, Nathanael Chambers, Niranjana Balasubramanian  
HIGHLIGHT: We propose a simple yet surprisingly effective strategy for improving event language models by perturbing event sequences so we can relax model dependence on text order.
- 648, TITLE: The Curse of Dense Low-Dimensional Information Retrieval for Large Index Sizes  
<https://aclanthology.org/2021.acl-short.77>  
AUTHORS: Nils Reimers, Iryna Gurevych  
HIGHLIGHT: We show theoretically and empirically that the performance for dense representations decreases quicker than sparse representations for increasing index sizes.
- 649, TITLE: Cross-lingual Text Classification with Heterogeneous Graph Neural Network  
<https://aclanthology.org/2021.acl-short.78>  
AUTHORS: Ziyun Wang, Xuan Liu, Peiji Yang, Shixing Liu, Zhisheng Wang  
HIGHLIGHT: In this paper we propose a simple yet effective method to incorporate heterogeneous information within and across languages for cross-lingual text classification using graph convolutional networks (GCN).
- 650, TITLE: Towards more equitable question answering systems: How much more data do you need?  
<https://aclanthology.org/2021.acl-short.79>  
AUTHORS: Arnab Debnath, Navid Rajabi, Fardina Fathmiul Alam, Antonios Anastasopoulos  
HIGHLIGHT: In this project we take a step back and study which approaches allow us to take the most advantage of existing resources in order to produce QA systems in many languages.
- 651, TITLE: Embedding Time Differences in Context-sensitive Neural Networks for Learning Time to Event  
<https://aclanthology.org/2021.acl-short.80>  
AUTHORS: Nazanin Dehghani, Hassan Hajipour, Hadi Amiri  
HIGHLIGHT: We propose an effective context-sensitive neural model for time to event (TTE) prediction task, which aims to predict the amount of time to/from the occurrence of given events in streaming content.
- 652, TITLE: Improving Compositional Generalization in Classification Tasks via Structure Annotations  
<https://aclanthology.org/2021.acl-short.81>  
AUTHORS: Juyong Kim, Pradeep Ravikumar, Joshua Ainslie, Santiago Ontanon  
HIGHLIGHT: In this work, we study compositional generalization in classification tasks and present two main contributions.
- 653, TITLE: Learning to Generate Task-Specific Adapters from Task Description  
<https://aclanthology.org/2021.acl-short.82>  
AUTHORS: Qinyuan Ye, Xiang Ren  
HIGHLIGHT: To this end, we introduce Hypter, a framework that improves text-to-text transformer's generalization ability to unseen tasks by training a hypernetwork to generate task-specific, light-weight adapters from task descriptions.
- 654, TITLE: QA-Driven Zero-shot Slot Filling with Weak Supervision Pretraining  
<https://aclanthology.org/2021.acl-short.83>  
AUTHORS: Xinya Du, Luheng He, Qi Li, Dian Yu, Panupong Pasupat, Yuan Zhang

**HIGHLIGHT:** In this work, we focus on the zero-shot slot-filling problem, where the model needs to predict slots and their values, given utterances from new domains without training on the target domain.

655, **TITLE:** Domain-Adaptive Pretraining Methods for Dialogue Understanding  
<https://aclanthology.org/2021.acl-short.84>  
**AUTHORS:** Han Wu, Kun Xu, Linfeng Song, Lifeng Jin, Haisong Zhang, Linqi Song  
**HIGHLIGHT:** In this paper, we probe the effectiveness of domain-adaptive pretraining objectives on downstream tasks.

656, **TITLE:** Targeting the Benchmark: On Methodology in Current Natural Language Processing Research  
<https://aclanthology.org/2021.acl-short.85>  
**AUTHORS:** David Schlangen  
**HIGHLIGHT:** In this paper, we try to step back for a moment from this pattern and work out possible argumentations and their parts.

657, **TITLE:** X-Fact: A New Benchmark Dataset for Multilingual Fact Checking  
<https://aclanthology.org/2021.acl-short.86>  
**AUTHORS:** Ashim Gupta, Vivek Srikumar  
**HIGHLIGHT:** In this work, we introduce : the largest publicly available multilingual dataset for factual verification of naturally existing real-world claims.

658, **TITLE:** nmT5 - Is parallel data still relevant for pre-training massively multilingual language models?  
<https://aclanthology.org/2021.acl-short.87>  
**AUTHORS:** Mihir Kale, Aditya Siddhant, Rami Al-Rfou, Linting Xue, Noah Constant, Melvin Johnson  
**HIGHLIGHT:** In this paper, we investigate the impact of incorporating parallel data into mT5 pre-training.

659, **TITLE:** Question Generation for Adaptive Education  
<https://aclanthology.org/2021.acl-short.88>  
**AUTHORS:** Megha Srivastava, Noah Goodman  
**HIGHLIGHT:** We explore targeted question generation as a controllable sequence generation task.

660, **TITLE:** A Simple Recipe for Multilingual Grammatical Error Correction  
<https://aclanthology.org/2021.acl-short.89>  
**AUTHORS:** Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, Aliaksei Severyn  
**HIGHLIGHT:** This paper presents a simple recipe to train state-of-the-art multilingual Grammatical Error Correction (GEC) models.

661, **TITLE:** Towards Visual Question Answering on Pathology Images  
<https://aclanthology.org/2021.acl-short.90>  
**AUTHORS:** Xuehai He, Zhuo Cai, Wenlan Wei, Yichen Zhang, Luntian Mou, Eric Xing, Pengtao Xie  
**HIGHLIGHT:** In this paper, we aim to develop a pathological visual question answering framework to analyze pathology images and answer medical questions related to these images.

662, **TITLE:** Efficient Text-based Reinforcement Learning by Jointly Leveraging State and Commonsense Graph Representations  
<https://aclanthology.org/2021.acl-short.91>  
**AUTHORS:** Keerthiram Murugesan, Mattia Atzeni, Pavan Kapanipathi, Kartik Talamadupula, Mrinmaya Sachan, Murray Campbell  
**HIGHLIGHT:** In this paper, we posit that to act efficiently in TBGs, an agent must be able to track the state of the game while retrieving and using relevant commonsense knowledge.

663, **TITLE:** mTVR: Multilingual Moment Retrieval in Videos  
<https://aclanthology.org/2021.acl-short.92>  
**AUTHORS:** Jie Lei, Tamara Berg, Mohit Bansal  
**HIGHLIGHT:** We introduce mTVR, a large-scale multilingual video moment retrieval dataset, containing 218K English and Chinese queries from 21.8K TV show video clips.

664, **TITLE:** Explicitly Capturing Relations between Entity Mentions via Graph Neural Networks for Domain-specific Named Entity Recognition  
<https://aclanthology.org/2021.acl-short.93>  
**AUTHORS:** Pei Chen, Haibo Ding, Jun Araki, Ruihong Huang

**HIGHLIGHT:** To address these challenges, we propose explicitly connecting entity mentions based on both global coreference relations and local dependency relations for building better entity mention representations.

665, **TITLE:** Improving Lexically Constrained Neural Machine Translation with Source-Conditioned Masked Span Prediction

<https://aclanthology.org/2021.acl-short.94>

**AUTHORS:** Gyubok Lee, Seongjun Yang, Edward Choi

**HIGHLIGHT:** In this paper, we instead tackle a more challenging setup consisting of domain-specific corpora with much longer n-gram and highly specialized terms.

666, **TITLE:** Quotation Recommendation and Interpretation Based on Transformation from Queries to Quotations

<https://aclanthology.org/2021.acl-short.95>

**AUTHORS:** Lingzhi Wang, Xingshan Zeng, Kam-Fai Wong

**HIGHLIGHT:** In this work, we introduce a transformation matrix that directly maps the query representations to quotation representations.

667, **TITLE:** Pre-training is a Hot Topic: Contextualized Document Embeddings Improve Topic Coherence

<https://aclanthology.org/2021.acl-short.96>

**AUTHORS:** Federico Bianchi, Silvia Terragni, Dirk Hovy

**HIGHLIGHT:** In this paper, we combine contextualized representations with neural topic models.

668, **TITLE:** Input Representations for Parsing Discourse Representation Structures: Comparing English with Chinese

<https://aclanthology.org/2021.acl-short.97>

**AUTHORS:** Chunliu Wang, Rik van Noord, Arianna Bisazza, Johan Bos

**HIGHLIGHT:** Overall, characters are preferred as input, both for English and Chinese.

669, **TITLE:** Code Generation from Natural Language with Less Prior Knowledge and More Monolingual Data

<https://aclanthology.org/2021.acl-short.98>

**AUTHORS:** Sajad Norouzi, Keyi Tang, Yanshuai Cao

**HIGHLIGHT:** This work investigates whether a generic transformer-based seq2seq model can achieve competitive performance with minimal code-generation-specific inductive bias design.

670, **TITLE:** Issues with Entailment-based Zero-shot Text Classification

<https://aclanthology.org/2021.acl-short.99>

**AUTHORS:** Tingting Ma, Jin-Ge Yao, Chin-Yew Lin, Tiejun Zhao

**HIGHLIGHT:** We observe huge variance across different classification datasets amongst standard BERT-based NLI models and surprisingly find that pre-trained BERT without any fine-tuning can yield competitive performance against BERT fine-tuned for NLI.

671, **TITLE:** Neural-Symbolic Commonsense Reasoner with Relation Predictors

<https://aclanthology.org/2021.acl-short.100>

**AUTHORS:** Farhad Moghimifar, Lizhen Qu, Terry Yue Zhuo, Gholamreza Haffari, Mahsa Baktashmotlagh

**HIGHLIGHT:** In this paper, we present a neural-symbolic reasoner, which is capable of reasoning over large-scale dynamic CKGs.

672, **TITLE:** What Motivates You? Benchmarking Automatic Detection of Basic Needs from Short Posts

<https://aclanthology.org/2021.acl-short.101>

**AUTHORS:** Sanja Stajner, Seren Yenikent, Bilal Ghanem, Marc Franco-Salvador

**HIGHLIGHT:** We benchmark the novel task of automatically detecting those needs on short posts in English, by modelling it as a ternary classification task, and as three binary classification tasks.

673, **TITLE:** Semantic Frame Induction using Masked Word Embeddings and Two-Step Clustering

<https://aclanthology.org/2021.acl-short.102>

**AUTHORS:** Kosuke Yamada, Ryohei Sasano, Koichi Takeda

**HIGHLIGHT:** To overcome these drawbacks, we propose a semantic frame induction method using masked word embeddings and two-step clustering.

674, **TITLE:** Lightweight Adapter Tuning for Multilingual Speech Translation

<https://aclanthology.org/2021.acl-short.103>

**AUTHORS:** Hang Le, Juan Pino, Changhan Wang, Jiatao Gu, Didier Schwab, Laurent Besacier

**HIGHLIGHT:** While adapter tuning was investigated for multilingual neural machine translation, this paper proposes a comprehensive analysis of adapters for multilingual speech translation (ST).

675, **TITLE:** Parameter Selection: Why We Should Pay More Attention to It  
<https://aclanthology.org/2021.acl-short.104>  
**AUTHORS:** Jie-Jyun Liu, Tsung-Han Yang, Si-An Chen, Chih-Jen Lin  
**HIGHLIGHT:** In this opinion paper, through an intriguing example we point out that the seriousness goes beyond what is generally recognized.

676, **TITLE:** Distinct Label Representations for Few-Shot Text Classification  
<https://aclanthology.org/2021.acl-short.105>  
**AUTHORS:** Sora Ohashi, Junya Takayama, Tomoyuki Kajiwara, Yuki Arase  
**HIGHLIGHT:** To address this problem, we propose a method that generates distinct label representations that embed information specific to each label.

677, **TITLE:** Learning to Solve NLP Tasks in an Incremental Number of Languages  
<https://aclanthology.org/2021.acl-short.106>  
**AUTHORS:** Giuseppe Castellucci, Simone Filice, Danilo Croce, Roberto Basili  
**HIGHLIGHT:** We propose a Continual Learning strategy that updates a model to support new languages over time, while maintaining consistent results on previously learned languages.

678, **TITLE:** Hi-Transformer: Hierarchical Interactive Transformer for Efficient and Effective Long Document Modeling  
<https://aclanthology.org/2021.acl-short.107>  
**AUTHORS:** Chuhan Wu, Fangzhao Wu, Tao Qi, Yongfeng Huang  
**HIGHLIGHT:** In order to handle this problem, we propose a hierarchical interactive Transformer (Hi-Transformer) for efficient and effective long document modeling.

679, **TITLE:** Robust Transfer Learning with Pretrained Language Models through Adapters  
<https://aclanthology.org/2021.acl-short.108>  
**AUTHORS:** Wenjuan Han, Bo Pang, Ying Nian Wu  
**HIGHLIGHT:** In particular, the performance considerably varies as the random seed changes or the number of pretraining and/or fine-tuning iterations varies, and the fine-tuned model is vulnerable to adversarial attack. We propose a simple yet effective adapter-based approach to mitigate these issues.

680, **TITLE:** Embracing Ambiguity: Shifting the Training Target of NLI Models  
<https://aclanthology.org/2021.acl-short.109>  
**AUTHORS:** Johannes Mario Meissner, Napat Thumwanit, Saku Sugawara, Akiko Aizawa  
**HIGHLIGHT:** In this paper, we explore the option of training directly on the estimated label distribution of the annotators in the NLI task, using a learning loss based on this ambiguity distribution instead of the gold-labels.

681, **TITLE:** Modeling Discriminative Representations for Out-of-Domain Detection with Supervised Contrastive Learning  
<https://aclanthology.org/2021.acl-short.110>  
**AUTHORS:** Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Zijun Liu, Yanan Wu, Hong Xu, Huixing Jiang, Weiran Xu  
**HIGHLIGHT:** In this paper, we propose a supervised contrastive learning objective to minimize intra-class variance by pulling together in-domain intents belonging to the same class and maximize inter-class variance by pushing apart samples from different classes.

682, **TITLE:** Preview, Attend and Review: Schema-Aware Curriculum Learning for Multi-Domain Dialogue State Tracking  
<https://aclanthology.org/2021.acl-short.111>  
**AUTHORS:** Yinpei Dai, Hangyu Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, Xiaodan Zhu  
**HIGHLIGHT:** In this paper, we propose to use curriculum learning (CL) to better leverage both the curriculum structure and schema structure for task-oriented dialogs.

683, **TITLE:** On the Generation of Medical Dialogs for COVID-19  
<https://aclanthology.org/2021.acl-short.112>  
**AUTHORS:** Meng Zhou, Zechen Li, Bowen Tan, Guangtao Zeng, Wenmian Yang, Xuehai He, Zeqian Ju, Subrato Chakravorty, Shu Chen, Xingyi Yang, Yichen Zhang, Qingyang Wu, Zhou Yu, Kun Xu, Eric Xing, Pengtao Xie  
**HIGHLIGHT:** To address this problem, we aim to develop a medical dialog system that can provide COVID19-related consultations.

684, TITLE: Constructing Multi-Modal Dialogue Dataset by Replacing Text with Semantically Relevant Images  
<https://aclanthology.org/2021.acl-short.113>  
AUTHORS: Nyoungwoo Lee, Suwon Shin, Jaegul Choo, Ho-Jin Choi, Sung-Hyon Myaeng  
HIGHLIGHT: In response, this paper proposes a 45k multi-modal dialogue dataset created with minimal human intervention.

685, TITLE: Exposing the limits of Zero-shot Cross-lingual Hate Speech Detection  
<https://aclanthology.org/2021.acl-short.114>  
AUTHORS: Debora Nozza  
HIGHLIGHT: This work is the first to shed light on the limits of this zero-shot, cross-lingual transfer learning framework for hate speech detection.

686, TITLE: BERTTune: Fine-Tuning Neural Machine Translation with BERTScore  
<https://aclanthology.org/2021.acl-short.115>  
AUTHORS: Inigo Jauregi Unanue, Jacob Parnell, Massimo Piccardi  
HIGHLIGHT: To amend this form of overfitting, in this paper we propose fine-tuning the models with a novel training objective based on the recently-proposed BERTScore evaluation metric.

687, TITLE: Entity Enhancement for Implicit Discourse Relation Classification in the Biomedical Domain  
<https://aclanthology.org/2021.acl-short.116>  
AUTHORS: Wei Shi, Vera Demberg  
HIGHLIGHT: We here tackle the task of implicit discourse relation classification on the biomedical domain, for which the Biomedical Discourse Relation Bank (BioDRB; Prasad et al., 2011) is available.

688, TITLE: Unsupervised Pronoun Resolution via Masked Noun-Phrase Prediction  
<https://aclanthology.org/2021.acl-short.117>  
AUTHORS: Ming Shen, Pratyay Banerjee, Chitta Baral  
HIGHLIGHT: In this work, we propose Masked Noun-Phrase Prediction (MNPP), a pre-training strategy to tackle pronoun resolution in a fully unsupervised setting.

689, TITLE: Addressing Semantic Drift in Generative Question Answering with Auxiliary Extraction  
<https://aclanthology.org/2021.acl-short.118>  
AUTHORS: Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang  
HIGHLIGHT: In this paper, we address these problems by a novel Rationale-Enriched Answer Generator (REAG), which incorporates an extractive mechanism into a generative model.

690, TITLE: Demoting the Lead Bias in News Summarization via Alternating Adversarial Learning  
<https://aclanthology.org/2021.acl-short.119>  
AUTHORS: Linzi Xing, Wen Xiao, Giuseppe Carenini  
HIGHLIGHT: In this paper, we introduce a novel technique to demote lead bias and make the summarizer focus more on the content semantics.

691, TITLE: DuReader\_robust: A Chinese Dataset Towards Evaluating Robustness and Generalization of Machine Reading Comprehension in Real-World Applications  
<https://aclanthology.org/2021.acl-short.120>  
AUTHORS: Hongxuan Tang, Hongyu Li, Jing Liu, Yu Hong, Hua Wu, Haifeng Wang  
HIGHLIGHT: In order to comprehensively verify the robustness and generalization of MRC models, we introduce a real-world Chinese dataset - DuReader\_robust .

692, TITLE: Sequence to General Tree: Knowledge-Guided Geometry Word Problem Solving  
<https://aclanthology.org/2021.acl-short.121>  
AUTHORS: Shih-hung Tsai, Chao-Chun Liang, Hsin-Min Wang, Keh-Yih Su  
HIGHLIGHT: In this paper, we propose sequence-to-general tree (S2G) that learns to generate interpretable and executable operation trees where the nodes can be formulas with an arbitrary number of arguments.

693, TITLE: Multi-Scale Progressive Attention Network for Video Question Answering  
<https://aclanthology.org/2021.acl-short.122>  
AUTHORS: Zhicheng Guo, Jiakuan Zhao, Licheng Jiao, Xu Liu, Lingling Li  
HIGHLIGHT: Therefore, we propose a novel Multi-Scale Progressive Attention Network (MSPAN) to achieve relational reasoning between cross-scale video information.

- 694, TITLE: Efficient Passage Retrieval with Hashing for Open-domain Question Answering  
<https://aclanthology.org/2021.acl-short.123>  
AUTHORS: Ikuya Yamada, Akari Asai, Hannaneh Hajishirzi  
HIGHLIGHT: In this paper, we introduce Binary Passage Retriever (BPR), a memory-efficient neural retrieval model that integrates a learning-to-hash technique into the state-of-the-art Dense Passage Retriever (DPR) to represent the passage index using compact binary codes rather than continuous vectors.
- 695, TITLE: Entity Concept-enhanced Few-shot Relation Extraction  
<https://aclanthology.org/2021.acl-short.124>  
AUTHORS: Shan Yang, Yongfei Zhang, Guanglin Niu, Qinghua Zhao, Shiliang Pu  
HIGHLIGHT: To address this problem, in this paper, we proposed a novel entity CONCEPT-enhanced FEw-shot Relation Extraction scheme (ConceptFERE), which introduces the inherent concepts of entities to provide clues for relation prediction and boost the relations classification performance.
- 696, TITLE: Improving Model Generalization: A Chinese Named Entity Recognition Case Study  
<https://aclanthology.org/2021.acl-short.125>  
AUTHORS: Guanqing Liang, Cane Wing-Ki Leung  
HIGHLIGHT: In this paper, we study the effect of data bias on model generalization, using Chinese Named Entity Recognition (NER) as a case study.
- 697, TITLE: Three Sentences Are All You Need: Local Path Enhanced Document Relation Extraction  
<https://aclanthology.org/2021.acl-short.126>  
AUTHORS: Quzhe Huang, Shengqi Zhu, Yansong Feng, Yuan Ye, Yuxuan Lai, Dongyan Zhao  
HIGHLIGHT: In this paper, we present an embarrassingly simple but effective method to heuristically select evidence sentences for document-level RE, which can be easily combined with BiLSTM to achieve good performance on benchmark datasets, even better than fancy graph neural network based methods.
- 698, TITLE: Unsupervised Cross-Domain Prerequisite Chain Learning using Variational Graph Autoencoders  
<https://aclanthology.org/2021.acl-short.127>  
AUTHORS: Irene Li, Vanessa Yan, Tianxiao Li, Rihao Qu, Dragomir Radev  
HIGHLIGHT: In this paper, we solve the task of unsupervised cross-domain concept prerequisite chain learning, using an optimized variational graph autoencoder.
- 699, TITLE: Attentive Multiview Text Representation for Differential Diagnosis  
<https://aclanthology.org/2021.acl-short.128>  
AUTHORS: Hadi Amiri, Mitra Mohtarami, Isaac Kohane  
HIGHLIGHT: We present a text representation approach that can combine different views (representations) of the same input through effective data fusion and attention strategies for ranking purposes.
- 700, TITLE: MedNLI Is Not Immune: Natural Language Inference Artifacts in the Clinical Domain  
<https://aclanthology.org/2021.acl-short.129>  
AUTHORS: Christine Herlihy, Rachel Rudinger  
HIGHLIGHT: We investigate whether MedNLI, a physician-annotated dataset with premises extracted from clinical notes, contains such artifacts (CITATION).
- 701, TITLE: Towards a more Robust Evaluation for Conversational Question Answering  
<https://aclanthology.org/2021.acl-short.130>  
AUTHORS: Wissam Sibli, Baris Sayil, Yacine Kessaci  
HIGHLIGHT: In this paper, we highlight this effect and propose new tools for evaluation and training in order to guard against the noted issues.
- 702, TITLE: VAULT: VARIable Unified Long Text Representation for Machine Reading Comprehension  
<https://aclanthology.org/2021.acl-short.131>  
AUTHORS: Haoyang Wen, Anthony Ferritto, Heng Ji, Radu Florian, Avi Sil  
HIGHLIGHT: In this work, we propose VAULT: a light-weight and parallel-efficient paragraph representation for MRC based on contextualized representation from long document input, trained using a new Gaussian distribution-based objective that pays close attention to the partially correct instances that are close to the ground-truth.
- 703, TITLE: Avoiding Overlap in Data Augmentation for AMR-to-Text Generation  
<https://aclanthology.org/2021.acl-short.132>  
AUTHORS: Wenchao Du, Jeffrey Flanigan



**HIGHLIGHT:** We propose methods for excluding parts of Gigaword to remove this overlap, and show that our approach leads to a more realistic evaluation of the task of AMR-to-text generation.

704, **TITLE:** Weakly-Supervised Methods for Suicide Risk Assessment: Role of Related Domains

<https://aclanthology.org/2021.acl-short.133>

**AUTHORS:** Chenghao Yang, Yudong Zhang, Smaranda Muresan

**HIGHLIGHT:** We propose an empirical investigation into several classes of weakly-supervised approaches, and show that using pseudo-labeling based on related issues around mental health (e.g., anxiety, depression) helps improve model performance for suicide risk assessment.

705, **TITLE:** Can Transformer Models Measure Coherence In Text: Re-Thinking the Shuffle Test

<https://aclanthology.org/2021.acl-short.134>

**AUTHORS:** Philippe Laban, Luke Dai, Lucas Bandarkar, Marti A. Hearst

**HIGHLIGHT:** We argue that this outstanding performance is unlikely to lead to a good model of text coherence, and suggest that the Shuffle Test should be approached in a Zero-Shot setting: models should be evaluated without being trained on the task itself.

706, **TITLE:** SimCLS: A Simple Framework for Contrastive Learning of Abstractive Summarization

<https://aclanthology.org/2021.acl-short.135>

**AUTHORS:** Yixin Liu, Pengfei Liu

**HIGHLIGHT:** In this paper, we present a conceptually simple while empirically powerful framework for abstractive summarization, SimCLS, which can bridge the gap between the learning objective and evaluation metrics resulting from the currently dominated sequence-to-sequence learning framework by formulating text generation as a reference-free evaluation problem (i.e., quality estimation) assisted by contrastive learning.

707, **TITLE:** SaRoCo: Detecting Satire in a Novel Romanian Corpus of News Articles

<https://aclanthology.org/2021.acl-short.136>

**AUTHORS:** Ana-Cristina Rogoz, Gaman Mihaela, Radu Tudor Ionescu

**HIGHLIGHT:** In this work, we introduce a corpus for satire detection in Romanian news.

708, **TITLE:** Bringing Structure into Summaries: a Faceted Summarization Dataset for Long Scientific Documents

<https://aclanthology.org/2021.acl-short.137>

**AUTHORS:** Rui Meng, Khushboo Thaker, Lei Zhang, Yue Dong, Xingdi Yuan, Tong Wang, Daqing He

**HIGHLIGHT:** In this study, we present FacetSum, a faceted summarization benchmark built on Emerald journal articles, covering a diverse range of domains.

709, **TITLE:** Replicating and Extending "Because Their Treebanks Leak": Graph Isomorphism, Covariants, and Parser Performance

<https://aclanthology.org/2021.acl-short.138>

**AUTHORS:** Mark Anderson, Anders S?gaard, Carlos G?mez-Rodr?guez

**HIGHLIGHT:** We present a replication study in which we also bin sentences by length and find that only a small subset of sentences vary in performance with respect to graph isomorphism.

710, **TITLE:** Don't Rule Out Monolingual Speakers: A Method For Crowdsourcing Machine Translation Data

<https://aclanthology.org/2021.acl-short.139>

**AUTHORS:** Rajat Bhatnagar, Ananya Ganesh, Katharina Kann

**HIGHLIGHT:** Here, we present a data collection strategy for MT which, in contrast, is cheap and simple, as it does not require bilingual speakers.