

- 1, TITLE: Social Norms Guide Reference Resolution
<https://aclanthology.org/2022.naacl-main.1>
AUTHORS: Mitchell Abrams, Matthias Scheutz
HIGHLIGHT: While some situated reference resolution is trivial, ambiguous cases arise when the language is underspecified or there are multiple candidate referents. This study investigates how pragmatic modulators external to the linguistic content are critical for the correct interpretation of referents in these scenarios.
- 2, TITLE: Learning Natural Language Generation with Truncated Reinforcement Learning
<https://aclanthology.org/2022.naacl-main.2>
AUTHORS: Alice Martin, Guillaume Quispe, Charles Ollion, Sylvain Le Corff, Florian Strub, Olivier Pietquin
HIGHLIGHT: This paper introduces TRuncated Reinforcement Learning for Language (TrufLL), an original approach to train conditional language models without a supervised learning phase, by only using reinforcement learning (RL).
- 3, TITLE: Language Model Augmented Monotonic Attention for Simultaneous Translation
<https://aclanthology.org/2022.naacl-main.3>
AUTHORS: Sathish Reddy Indurthi, Mohd Abbas Zaidi, Beomseok Lee, Nikhil Kumar Lakumarapu, Sangha Kim
HIGHLIGHT: In this work, we propose a framework to aid monotonic attention with an external language model to improve its decisions.
- 4, TITLE: What Makes a Good and Useful Summary? Incorporating Users in Automatic Summarization Research
<https://aclanthology.org/2022.naacl-main.4>
AUTHORS: Maartje Ter Hoeve, Julia Kiseleva, Maarten Rijke
HIGHLIGHT: In this work we focus on university students, who make extensive use of summaries during their studies.
- 5, TITLE: ErAConD: Error Annotated Conversational Dialog Dataset for Grammatical Error Correction
<https://aclanthology.org/2022.naacl-main.5>
AUTHORS: Xun Yuan, Derek Pham, Sam Davidson, Zhou Yu
HIGHLIGHT: In this paper, we present a novel GEC dataset consisting of parallel original and corrected utterances drawn from open-domain chatbot conversations; this dataset is, to our knowledge, the first GEC dataset targeted to a human-machine conversational setting.
- 6, TITLE: Semantic Diversity in Dialogue with Natural Language Inference
<https://aclanthology.org/2022.naacl-main.6>
AUTHORS: Katherine Stasaski, Marti Hearst
HIGHLIGHT: This paper makes two substantial contributions to improving diversity in dialogue generation.
- 7, TITLE: LEA: Meta Knowledge-Driven Self-Attentive Document Embedding for Few-Shot Text Classification
<https://aclanthology.org/2022.naacl-main.7>
AUTHORS: S. K. Hong, Tae Young Jang
HIGHLIGHT: In the study, we propose a novel learning method for learning how to attend, called LEA, through which meta-level attention aspects are derived based on our meta-learning strategy.
- 8, TITLE: Enhancing Self-Attention with Knowledge-Assisted Attention Maps
<https://aclanthology.org/2022.naacl-main.8>
AUTHORS: Jiangang Bai, Yujing Wang, Hong Sun, Ruonan Wu, Tianmeng Yang, Pengfei Tang, Defu Cao, Mingliang Zhang, Yunhai Tong, Yaming Yang, Jing Bai, Ruofei Zhang, Hao Sun, Wei Shen
HIGHLIGHT: In this paper, we propose a novel and generic solution, KAM-BERT, which directly incorporates knowledge-generated attention maps into the self-attention mechanism.
- 9, TITLE: Batch-Softmax Contrastive Loss for Pairwise Sentence Scoring Tasks
<https://aclanthology.org/2022.naacl-main.9>
AUTHORS: Anton Chernyavskiy, Dmitry Ilvovsky, Pavel Kalinin, Preslav Nakov
HIGHLIGHT: We introduce and study a number of variations in the calculation of the loss as well as in the overall training procedure; in particular, we find that a special data shuffling can be quite important.
- 10, TITLE: NewsEdits: A News Article Revision Dataset and a Novel Document-Level Reasoning Challenge
<https://aclanthology.org/2022.naacl-main.10>
AUTHORS: Alexander Spangher, Xiang Ren, Jonathan May, Nanyun Peng
HIGHLIGHT: News article revision histories provide clues to narrative and factual evolution in news articles. To facilitate analysis of this evolution, we present the first publicly available dataset of news revision histories, NewsEdits.
- 11, TITLE: Putting the Con in Context: Identifying Deceptive Actors in the Game of Mafia
<https://aclanthology.org/2022.naacl-main.11>
AUTHORS: Samee Ibraheem, Gaoyue Zhou, John DeNero
HIGHLIGHT: In this work, we analyze the effect of speaker role on language use through the game of Mafia, in which participants are assigned either an honest or a deceptive role.
- 12, TITLE: SUBS: Subtree Substitution for Compositional Semantic Parsing
<https://aclanthology.org/2022.naacl-main.12>
AUTHORS: Jingfeng Yang, Le Zhang, Diyi Yang

HIGHLIGHT: We propose to use subtree substitution for compositional data augmentation, where we consider subtrees with similar semantic functions as exchangeable.

13, **TITLE:** Two Contrasting Data Annotation Paradigms for Subjective NLP Tasks

<https://aclanthology.org/2022.naacl-main.13>

AUTHORS: Paul Rottger, Bertie Vidgen, Dirk Hovy, Janet Pierrehumbert

HIGHLIGHT: This has led to partly-subjective datasets that fail to serve a clear downstream use. To address this issue, we propose two contrasting paradigms for data annotation.

14, **TITLE:** Do Deep Neural Nets Display Human-like Attention in Short Answer Scoring?

<https://aclanthology.org/2022.naacl-main.14>

AUTHORS: Zijie Zeng, Xinyu Li, Dragan Gasevic, Guanliang Chen

HIGHLIGHT: This study aimed to investigate whether (and to what extent) DL-based graders align with human graders regarding the important words they identify when marking short answer questions.

15, **TITLE:** Knowledge-Grounded Dialogue Generation with a Unified Knowledge Representation

<https://aclanthology.org/2022.naacl-main.15>

AUTHORS: Yu Li, Baolin Peng, Yelong Shen, Yi Mao, Lars Liden, Zhou Yu, Jianfeng Gao

HIGHLIGHT: In addition, it is challenging to generalize to the domains that require different types of knowledge sources. To address the above challenges, we present PLUG, a language model that homogenizes different knowledge sources to a unified knowledge representation for knowledge-grounded dialogue generation tasks.

16, **TITLE:** CERES: Pretraining of Graph-Conditioned Transformer for Semi-Structured Session Data

<https://aclanthology.org/2022.naacl-main.16>

AUTHORS: Rui Feng, Chen Luo, Qingyu Yin, Bing Yin, Tuo Zhao, Chao Zhang

HIGHLIGHT: Despite recent advances in self-supervised learning for text or graphs, there lack of self-supervised learning models that can effectively capture both intra-item semantics and inter-item interactions for semi-structured sessions. To fill this gap, we propose CERES, a graph-based transformer model for semi-structured session data.

17, **TITLE:** Political Ideology and Polarization: A Multi-dimensional Approach

<https://aclanthology.org/2022.naacl-main.17>

AUTHORS: Barea Sinno, Bernardo Oviedo, Katherine Atwell, Malihe Alikhani, Junyi Jessy Li

HIGHLIGHT: Recent research has made great strides towards understanding the ideological bias (i.e., stance) of news media along the left-right spectrum. In this work, we instead take a novel and more nuanced approach for the study of ideology based on its left or right positions on the issue being discussed.

18, **TITLE:** Cooperative Self-training of Machine Reading Comprehension

<https://aclanthology.org/2022.naacl-main.18>

AUTHORS: Hongyin Luo, Shang-Wen Li, Mingye Gao, Seunghak Yu, James Glass

HIGHLIGHT: In this work, we propose a cooperative self-training framework, RGX, for automatically generating more non-trivial question-answer pairs to improve model performance.

19, **TITLE:** GlobEnc: Quantifying Global Token Attribution by Incorporating the Whole Encoder Layer in Transformers

<https://aclanthology.org/2022.naacl-main.19>

AUTHORS: Ali Modarressi, Mohsen Fayyaz, Yadollah Yaghoobzadeh, Mohammad Taher Pilehvar

HIGHLIGHT: This paper introduces a novel token attribution analysis method that incorporates all the components in the encoder block and aggregates this throughout layers.

20, **TITLE:** A Robustly Optimized BMRC for Aspect Sentiment Triplet Extraction

<https://aclanthology.org/2022.naacl-main.20>

AUTHORS: Shu Liu, Kaiwen Li, Zuhe Li

HIGHLIGHT: The bidirectional machine reading comprehension (BMRC), can effectively deal with ASTE task, but several problems remains, such as query conflict and probability unilateral decrease. Therefore, this paper presents a robustly optimized BMRC method by incorporating four improvements.

21, **TITLE:** Seed-Guided Topic Discovery with Out-of-Vocabulary Seeds

<https://aclanthology.org/2022.naacl-main.21>

AUTHORS: Yu Zhang, Yu Meng, Xuan Wang, Sheng Wang, Jiawei Han

HIGHLIGHT: In this paper, we generalize the task of seed-guided topic discovery to allow out-of-vocabulary seeds.

22, **TITLE:** Towards Process-Oriented, Modular, and Versatile Question Generation that Meets Educational Needs

<https://aclanthology.org/2022.naacl-main.22>

AUTHORS: Xu Wang, Simin Fan, Jessica Houghton, Lu Wang

HIGHLIGHT: In this work, we aim to pinpoint key impediments and investigate how to improve the usability of automatic QG techniques for educational purposes by understanding how instructors construct questions and identifying touch points to enhance the underlying NLP models.

23, **TITLE:** SwahBERT: Language Model of Swahili

<https://aclanthology.org/2022.naacl-main.23>

AUTHORS: Gati Martin, Medard Edmund Mswahili, Young-Seob Jeong, Jeong Young-Seob

HIGHLIGHT: Thanks to the recent growth of online forums and news platforms of Swahili, we introduce two datasets of Swahili in this paper: a pre-training dataset of approximately 105MB with 16M words and annotated dataset of 13K instances for the emotion classification task.

24, **TITLE:** Deconstructing NLG Evaluation: Evaluation Practices, Assumptions, and Their Implications
<https://aclanthology.org/2022.naacl-main.24>
AUTHORS: Kaitlyn Zhou, Su Lin Blodgett, Adam Trischler, Hal Daum? III, Kaheer Suleman, Alexandra Olteanu
HIGHLIGHT: Combining a formative semi-structured interview study of NLG practitioners (N=18) with a survey study of a broader sample of practitioners (N=61), we surface goals, community practices, assumptions, and constraints that shape NLG evaluations, examining their implications and how they embody ethical considerations.

25, **TITLE:** TSTR: Too Short to Represent, Summarize with Details! Intro-Guided Extended Summary Generation
<https://aclanthology.org/2022.naacl-main.25>
AUTHORS: Sajad Sotudeh, Nazli Goharian
HIGHLIGHT: In this paper, we propose TSTR, an extractive summarizer that utilizes the introductory information of documents as pointers to their salient information.

26, **TITLE:** Empathic Machines: Using Intermediate Features as Levers to Emulate Emotions in Text-To-Speech Systems
<https://aclanthology.org/2022.naacl-main.26>
AUTHORS: Saiteja Kosgi, Sarath Sivaprasad, Niranjan Pedanekar, Anil Nelakanti, Vineet Gandhi
HIGHLIGHT: We present a method to control the emotional prosody of Text to Speech (TTS) systems by using phoneme-level intermediate features (pitch, energy, and duration) as levers.

27, **TITLE:** The Why and The How: A Survey on Natural Language Interaction in Visualization
<https://aclanthology.org/2022.naacl-main.27>
AUTHORS: Henrik Voigt, Ozge Alacam, Monique Meuschke, Kai Lawonn, Sina Zarrie?
HIGHLIGHT: In this survey, we provide an overview of natural language-based interaction in the research area of visualization.

28, **TITLE:** Understand before Answer: Improve Temporal Reading Comprehension via Precise Question Understanding
<https://aclanthology.org/2022.naacl-main.28>
AUTHORS: Hao Huang, Xiubo Geng, Guodong Long, Daxin Jiang
HIGHLIGHT: This work studies temporal reading comprehension (TRC), which reads a free-text passage and answers temporal ordering questions.

29, **TITLE:** User-Driven Research of Medical Note Generation Software
<https://aclanthology.org/2022.naacl-main.29>
AUTHORS: Tom Knoll, Francesco Moramarco, Alex Papadopoulos Korfiatis, Rachel Young, Claudia Ruffini, Mark Perera, Christian Perstl, Ehud Reiter, Anya Belz, Aleksandar Savkov
HIGHLIGHT: In this paper, we present three rounds of user studies, carried out in the context of developing a medical note generation system.

30, **TITLE:** Ask Me Anything in Your Native Language
<https://aclanthology.org/2022.naacl-main.30>
AUTHORS: Nikita Sorokin, Dmitry Abulkhanov, Irina Piontkovskaya, Valentin Malykh
HIGHLIGHT: We present a novel approach based on single encoder for query and passage for retrieval from multi-lingual collection, together with cross-lingual generative reader.

31, **TITLE:** Diversifying Neural Dialogue Generation via Negative Distillation
<https://aclanthology.org/2022.naacl-main.31>
AUTHORS: Yiwei Li, Shaoxiong Feng, Bin Sun, Kan Li
HIGHLIGHT: However, its performance is hindered by two issues, ignoring low-frequency but generic responses and bringing low-frequency but meaningless responses. In this paper, we propose a novel negative training paradigm, called negative distillation, to keep the model away from the undesirable generic responses while avoiding the above problems.

32, **TITLE:** On Synthetic Data for Back Translation
<https://aclanthology.org/2022.naacl-main.32>
AUTHORS: Jiahao Xu, Yubin Ruan, Wei Bi, Guoping Huang, Shuming Shi, Lihui Chen, Lemao Liu
HIGHLIGHT: Through both theoretical and empirical studies, we identify two key factors on synthetic data controlling the back-translation NMT performance, which are quality and importance.

33, **TITLE:** Mapping the Design Space of Human-AI Interaction in Text Summarization
<https://aclanthology.org/2022.naacl-main.33>
AUTHORS: Ruijia Cheng, Alison Smith-Renner, Ke Zhang, Joel Tetreault, Alejandro Jaimes-Larrarte
HIGHLIGHT: We first conducted a systematic literature review of 70 papers, developing a taxonomy of five interactions in AI-assisted text generation and relevant design dimensions. We designed text summarization prototypes for each interaction.

34, **TITLE:** Towards Robust and Semantically Organised Latent Representations for Unsupervised Text Style Transfer
<https://aclanthology.org/2022.naacl-main.34>
AUTHORS: Sharan Narasimhan, Suvodip Dey, Maunendra Desarkar

HIGHLIGHT: We introduce EPAAEs (Embedding Perturbed Adversarial AutoEncoders) which completes this perturbation model, by adding a finely adjustable noise component on the continuous embeddings space.

35, **TITLE:** An Exploration of Post-Editing Effectiveness in Text Summarization
<https://aclanthology.org/2022.naacl-main.35>

AUTHORS: Vivian Lai, Alison Smith-Renner, Ke Zhang, Ruijia Cheng, Wenjuan Zhang, Joel Tetreault, Alejandro Jaimes-Larrarte

HIGHLIGHT: Therefore, we explored whether post-editing offers advantages in text summarization.

36, **TITLE:** Automatic Correction of Human Translations
<https://aclanthology.org/2022.naacl-main.36>

AUTHORS: Jessy Lin, Geza Kovacs, Aditya Shastri, Joern Wuebker, John DeNero

HIGHLIGHT: We introduce translation error correction (TEC), the task of automatically correcting human-generated translations. Imperfections in machine translations (MT) have long motivated systems for improving translations post-hoc with automatic post-editing.

37, **TITLE:** On the Robustness of Reading Comprehension Models to Entity Renaming
<https://aclanthology.org/2022.naacl-main.37>

AUTHORS: Jun Yan, Yang Xiao, Sagnik Mukherjee, Bill Yuchen Lin, Robin Jia, Xiang Ren

HIGHLIGHT: Such failures imply that models overly rely on entity information to answer questions, and thus may generalize poorly when facts about the world change or questions are asked about novel entities. To systematically audit this issue, we present a pipeline to automatically generate test examples at scale, by replacing entity names in the original test sample with names from a variety of sources, ranging from names in the same test set, to common names in life, to arbitrary strings.

38, **TITLE:** Explaining Why: How Instructions and User Interfaces Impact Annotator Rationales When Labeling Text Data
<https://aclanthology.org/2022.naacl-main.38>

AUTHORS: Cynthia Sullivan, William Brackenbury, Andrew McNut, Kevin Bryson, Kbyllofficial@gmail.com
Kbyllofficial@gmail.com, Yuxin Chen, Michael Littman, Chenhao Tan, Blase Ur

HIGHLIGHT: We conducted a 332-participant online user study to understand how humans select rationales, especially how different instructions and user interface affordances impact the rationales chosen.

39, **TITLE:** Fine-tuning Pre-trained Language Models for Few-shot Intent Detection: Supervised Pre-training and Isotropization
<https://aclanthology.org/2022.naacl-main.39>

AUTHORS: Haode Zhang, Haowen Liang, Yuwei Zhang, Li-Ming Zhan, Xiao-Ming Wu, Xiaolei Lu, Albert Lam

HIGHLIGHT: Inspired by recent research in isotropization, we propose to improve supervised pre-training by regularizing the feature space towards isotropy.

40, **TITLE:** Cross-document Misinformation Detection based on Event Graph Reasoning
<https://aclanthology.org/2022.naacl-main.40>

AUTHORS: Xueqing Wu, Kung-Hsiang Huang, Yi Fung, Heng Ji

HIGHLIGHT: Multiple news articles may contain complementary or contradictory information that readers can leverage to help detect fake news. Inspired by this process, we propose a novel task of cross-document misinformation detection.

41, **TITLE:** Disentangled Action Recognition with Knowledge Bases
<https://aclanthology.org/2022.naacl-main.41>

AUTHORS: Zhekun Luo, Shalini Ghosh, Devin Guillory, Keizo Kato, Trevor Darrell, Huijuan Xu

HIGHLIGHT: In this paper, we aim to improve the generalization ability of the compositional action recognition model to novel verbs or novel nouns that are unseen during training time, by leveraging the power of knowledge graphs.

42, **TITLE:** Machine-in-the-Loop Rewriting for Creative Image Captioning
<https://aclanthology.org/2022.naacl-main.42>

AUTHORS: Vishakh Padmakumar, He He

HIGHLIGHT: To allow the user to retain control over the content, we train a rewriting model that, when prompted, modifies specified spans of text within the user’s original draft to introduce descriptive and figurative elements in the text.

43, **TITLE:** A Word is Worth A Thousand Dollars: Adversarial Attack on Tweets Fools Stock Prediction
<https://aclanthology.org/2022.naacl-main.43>

AUTHORS: Yong Xie, Dakuo Wang, Pin-Yu Chen, Jinjun Xiong, Sijia Liu, Oluwasanmi Koyejo

HIGHLIGHT: In this paper, we experiment with a variety of adversarial attack configurations to fool three stock prediction victim models.

44, **TITLE:** Building Multilingual Machine Translation Systems That Serve Arbitrary XY Translations
<https://aclanthology.org/2022.naacl-main.44>

AUTHORS: Akiko Eriguchi, Shufang Xie, Tao Qin, Hany Hassan

HIGHLIGHT: The model suffers from poor performance in one-to-many and many-to-many with zero-shot setup. To address this issue, this paper discusses how to practically build MNMT systems that serve arbitrary X-Y translation directions while leveraging multilinguality with a two-stage training strategy of pretraining and finetuning.

- 45, TITLE: Non-Autoregressive Neural Machine Translation with Consistency Regularization Optimized Variational Framework
<https://aclanthology.org/2022.naacl-main.45>
AUTHORS: Minghao Zhu, Junli Wang, Chungang Yan
HIGHLIGHT: One of the prominent VAE-based NAT frameworks, LaNMT, achieves great improvements to vanilla models, but still suffers from two main issues which lower down the translation quality: (1) mismatch between training and inference circumstances and (2) inadequacy of latent representations. In this work, we target on addressing these issues by proposing posterior consistency regularization.
- 46, TITLE: User-Centric Gender Rewriting
<https://aclanthology.org/2022.naacl-main.46>
AUTHORS: Bashar Alhafni, Nizar Habash, Houda Bouamor
HIGHLIGHT: In this paper, we define the task of gender rewriting in contexts involving two users (I and/or You) - first and second grammatical persons with independent grammatical gender preferences.
- 47, TITLE: Reframing Human-AI Collaboration for Generating Free-Text Explanations
<https://aclanthology.org/2022.naacl-main.47>
AUTHORS: Sarah Wiegreffe, Jack Hessel, Swabha Swayamdipta, Mark Riedl, Yejin Choi
HIGHLIGHT: We consider the task of generating free-text explanations using human-written examples in a few-shot manner.
- 48, TITLE: EmRel: Joint Representation of Entities and Embedded Relations for Multi-triple Extraction
<https://aclanthology.org/2022.naacl-main.48>
AUTHORS: Benfeng Xu, Quan Wang, Yajuan Lyu, Yabing Shi, Yong Zhu, Jie Gao, Zhendong Mao
HIGHLIGHT: While existing works only explore entity representations, we propose to explicitly introduce relation representation, jointly represent it with entities, and novelly align them to identify valid triples.
- 49, TITLE: Meta Learning for Natural Language Processing: A Survey
<https://aclanthology.org/2022.naacl-main.49>
AUTHORS: Hung-yi Lee, Shang-Wen Li, Thang Vu
HIGHLIGHT: Our goal with this survey paper is to offer researchers pointers to relevant meta-learning works in NLP and attract more attention from the NLP community to drive future innovation.
- 50, TITLE: Analyzing Modality Robustness in Multimodal Sentiment Analysis
<https://aclanthology.org/2022.naacl-main.50>
AUTHORS: Devamanyu Hazarika, Yingting Li, Bo Cheng, Shuai Zhao, Roger Zimmermann, Soujanya Poria
HIGHLIGHT: In this work, we hope to address that by (i) Proposing simple diagnostic checks for modality robustness in a trained multimodal model.
- 51, TITLE: Fuse It More Deeply! A Variational Transformer with Layer-Wise Latent Variable Inference for Text Generation
<https://aclanthology.org/2022.naacl-main.51>
AUTHORS: Jinyi Hu, Xiaoyuan Yi, Wenhao Li, Maosong Sun, Xing Xie
HIGHLIGHT: However, due to the sequential nature of the text, auto-regressive decoders tend to ignore latent variables and then reduce to simple language models, known as the KL vanishing problem, which would further deteriorate when VAE is combined with Transformer-based structures. To ameliorate this problem, we propose Della, a novel variational Transformer framework.
- 52, TITLE: Easy Adaptation to Mitigate Gender Bias in Multilingual Text Classification
<https://aclanthology.org/2022.naacl-main.52>
AUTHORS: Xiaolei Huang
HIGHLIGHT: In this work, we treat the gender as domains (e.g., male vs. female) and present a standard domain adaptation model to reduce the gender bias and improve performance of text classifiers under multilingual settings.
- 53, TITLE: On the Use of External Data for Spoken Named Entity Recognition
<https://aclanthology.org/2022.naacl-main.53>
AUTHORS: Ankita Pasad, Felix Wu, Suwon Shon, Karen Livescu, Kyu Han
HIGHLIGHT: In this work, we focus on low-resource spoken named entity recognition (NER) and address the question: Beyond self-supervised pre-training, how can we use external speech and/or text data that are not annotated for the task?
- 54, TITLE: Long-term Control for Dialogue Generation: Methods and Evaluation
<https://aclanthology.org/2022.naacl-main.54>
AUTHORS: Ramya Ramakrishnan, Hashan Narangodage, Mauro Schilman, Kilian Weinberger, Ryan McDonald
HIGHLIGHT: In this work, we focus on constrained long-term dialogue generation, which involves more fine-grained control and requires a given set of control words to appear in generated responses.
- 55, TITLE: Learning Dialogue Representations from Consecutive Utterances
<https://aclanthology.org/2022.naacl-main.55>
AUTHORS: Zhihan Zhou, Dejiao Zhang, Wei Xiao, Nicholas Dingwall, Xiaofei Ma, Andrew Arnold, Bing Xiang
HIGHLIGHT: In this paper, we introduce Dialogue Sentence Embedding (DSE), a self-supervised contrastive learning method that learns effective dialogue representations suitable for a wide range of dialogue tasks.

- 56, TITLE: On the Machine Learning of Ethical Judgments from Natural Language
<https://aclanthology.org/2022.naacl-main.56>
AUTHORS: Zeerak Talat, Hagen Blix, Josef Valvoda, Maya Indira Ganesh, Ryan Cotterell, Adina Williams
HIGHLIGHT: One recent approach in this vein is the construction of NLP morality models that can take in arbitrary text and output a moral judgment about the situation described. In this work, we offer a critique of such NLP methods for automating ethical decision-making.
- 57, TITLE: NeuroLogic A*esque Decoding: Constrained Text Generation with Lookahead Heuristics
<https://aclanthology.org/2022.naacl-main.57>
AUTHORS: Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah Smith, Yejin Choi
HIGHLIGHT: Drawing inspiration from the A* search algorithm, we propose NeuroLogic A*esque, a decoding algorithm that incorporates heuristic estimates of future cost.
- 58, TITLE: PARADISE: Exploiting Parallel Data for Multilingual Sequence-to-Sequence Pretraining
<https://aclanthology.org/2022.naacl-main.58>
AUTHORS: Machel Reid, Mikel Artetxe
HIGHLIGHT: In this paper, we present PARADISE (PARAllel &Denoising Integration in SEquence-to-sequence models), which extends the conventional denoising objective used to train these models by (i) replacing words in the noised sequence according to a multilingual dictionary, and (ii) predicting the reference translation according to a parallel corpus instead of recovering the original sequence.
- 59, TITLE: Explaining Toxic Text via Knowledge Enhanced Text Generation
<https://aclanthology.org/2022.naacl-main.59>
AUTHORS: Rohit Sridhar, Diyi Yang
HIGHLIGHT: Previous literature has mostly focused on classifying and detecting toxic speech, and existing efforts on explaining stereotypes in toxic speech mainly use standard text generation approaches, resulting in generic and repetitive explanations. Building on these prior works, we introduce a novel knowledge-informed encoder-decoder framework to utilize multiple knowledge sources to generate implications of biased text.
- 60, TITLE: Teaching BERT to Wait: Balancing Accuracy and Latency for Streaming Disfluency Detection
<https://aclanthology.org/2022.naacl-main.60>
AUTHORS: Angelica Chen, Vicky Zayats, Daniel Walker, Dirk Padfield
HIGHLIGHT: In this work we propose a streaming BERT-based sequence tagging model that, combined with a novel training objective, is capable of detecting disfluencies in real-time while balancing accuracy and latency.
- 61, TITLE: GRAM: Fast Fine-tuning of Pre-trained Language Models for Content-based Collaborative Filtering
<https://aclanthology.org/2022.naacl-main.61>
AUTHORS: Yoonseok Yang, Kyu Seok Kim, Minsam Kim, Juneyoung Park
HIGHLIGHT: However, it is resource-intensive to train a PLM-based CCF model in an end-to-end (E2E) manner, since optimization involves back-propagating through every content encoding within a given user interaction sequence. To tackle this issue, we propose GRAM (GRAdient Accumulation for Multi-modality in CCF), which exploits the fact that a given item often appears multiple times within a batch of interaction histories.
- 62, TITLE: Generating Repetitions with Appropriate Repeated Words
<https://aclanthology.org/2022.naacl-main.62>
AUTHORS: Toshiki Kawamoto, Hidetaka Kamigaito, Kotaro Funakoshi, Manabu Okumura
HIGHLIGHT: In this work, we focus on repetition generation.
- 63, TITLE: Textless Speech-to-Speech Translation on Real Data
<https://aclanthology.org/2022.naacl-main.63>
AUTHORS: Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Yossi Adi, Juan Pino, Jiatao Gu, Wei-Ning Hsu
HIGHLIGHT: We present a textless speech-to-speech translation (S2ST) system that can translate speech from one language into another language and can be built without the need of any text data.
- 64, TITLE: WALNUT: A Benchmark on Semi-weakly Supervised Learning for Natural Language Understanding
<https://aclanthology.org/2022.naacl-main.64>
AUTHORS: Guoqing Zheng, Giannis Karamanolakis, Kai Shu, Ahmed Awadallah
HIGHLIGHT: It is thus hard to compare different approaches and evaluate the benefit of weak supervision without access to a unified and systematic benchmark with diverse tasks and real-world weak labeling rules. In this paper, we propose such a benchmark, named WALNUT, to advocate and facilitate research on weak supervision for NLU.
- 65, TITLE: CompactIE: Compact Facts in Open Information Extraction
<https://aclanthology.org/2022.naacl-main.65>
AUTHORS: Farima Fatahi Bayat, Nikita Bhutani, H. Jagadish
HIGHLIGHT: We propose CompactIE, an OpenIE system that uses a novel pipelined approach to produce compact extractions with overlapping constituents.

- 66, TITLE: CoSIm: Commonsense Reasoning for Counterfactual Scene Imagination
<https://aclanthology.org/2022.naacl-main.66>
AUTHORS: Hyounghun Kim, Abhay Zala, Mohit Bansal
HIGHLIGHT: We present a baseline model based on a vision-language Transformer (i.e., LXMERT) and ablation studies.
- 67, TITLE: Abstraction not Memory: BERT and the English Article System
<https://aclanthology.org/2022.naacl-main.67>
AUTHORS: Harish Tayyar Madabushi, Dagmar Divjak, Petar Milin
HIGHLIGHT: To this end, we compare the performance of native English speakers and pre-trained models on the task of article prediction set up as a three way choice (a/an, the, zero).
- 68, TITLE: OmniTab: Pretraining with Natural and Synthetic Data for Few-shot Table-based Question Answering
<https://aclanthology.org/2022.naacl-main.68>
AUTHORS: Zhengbao Jiang, Yi Mao, Pengcheng He, Graham Neubig, Weizhu Chen
HIGHLIGHT: Motivated by the fact that table-based QA requires both alignment between questions and tables and the ability to perform complicated reasoning over multiple table elements, we propose an omnivorous pretraining approach that consumes both natural and synthetic data to endow models with these respective abilities.
- 69, TITLE: Provably Confidential Language Modelling
<https://aclanthology.org/2022.naacl-main.69>
AUTHORS: Xuandong Zhao, Lei Li, Yu-Xiang Wang
HIGHLIGHT: In this paper, we propose Confidentially Redacted Training (CRT), a method to train language generation models while protecting the confidential segments.
- 70, TITLE: KAT: A Knowledge Augmented Transformer for Vision-and-Language
<https://aclanthology.org/2022.naacl-main.70>
AUTHORS: Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander Hauptmann, Yonatan Bisk, Jianfeng Gao
HIGHLIGHT: In this work, we ask a complementary question: Can multimodal transformers leverage explicit knowledge in their reasoning?
- 71, TITLE: When a sentence does not introduce a discourse entity, Transformer-based models still sometimes refer to it
<https://aclanthology.org/2022.naacl-main.71>
AUTHORS: Sebastian Schuster, Tal Linzen
HIGHLIGHT: In this work, we adapt the psycholinguistic assessment of language models paradigm to higher-level linguistic phenomena and introduce an English evaluation suite that targets the knowledge of the interactions between sentential operators and indefinite NPs.
- 72, TITLE: On Curriculum Learning for Commonsense Reasoning
<https://aclanthology.org/2022.naacl-main.72>
AUTHORS: Adyasha Maharana, Mohit Bansal
HIGHLIGHT: We use paced curriculum learning to rank data and sample training mini-batches with increasing levels of difficulty from the ranked dataset during finetuning. Further, we investigate the effect of an adaptive curriculum, i.e., the data ranking is dynamically updated during training based on the current state of the learner model.
- 73, TITLE: DocTime: A Document-level Temporal Dependency Graph Parser
<https://aclanthology.org/2022.naacl-main.73>
AUTHORS: Puneet Mathur, Vlad Morariu, Verena Kaynig-Fittkau, Jiuxiang Gu, Franck Dernoncourt, Quan Tran, Ani Nenkova, Dinesh Manocha, Rajiv Jain
HIGHLIGHT: We introduce DocTime - a novel temporal dependency graph (TDG) parser that takes as input a text document and produces a temporal dependency graph.
- 74, TITLE: FactPEGASUS: Factuality-Aware Pre-training and Fine-tuning for Abstractive Summarization
<https://aclanthology.org/2022.naacl-main.74>
AUTHORS: David Wan, Mohit Bansal
HIGHLIGHT: We present FactPEGASUS, an abstractive summarization model that addresses the problem of factuality during pre-training and fine-tuning: (1) We augment the sentence selection strategy of PEGASUS's (Zhang et al., 2019) pre-training objective to create pseudo-summaries that are both important and factual; (2) We introduce three complementary components for fine-tuning.
- 75, TITLE: ScAN: Suicide Attempt and Ideation Events Dataset
<https://aclanthology.org/2022.naacl-main.75>
AUTHORS: Bhanu Pratap Singh Rawat, Samuel Kovaly, Hong Yu, Wilfred Pigeon
HIGHLIGHT: In this study, we first built Suicide Attempt and Ideation Events (ScAN) dataset, a subset of the publicly available MIMIC III dataset spanning over 12k+ EHR notes with 19k+ annotated SA and SI events information.
- 76, TITLE: Socially Aware Bias Measurements for Hindi Language Representations
<https://aclanthology.org/2022.naacl-main.76>
AUTHORS: Vijit Malik, Sunipa Dev, Akihiro Nishi, Nanyun Peng, Kai-Wei Chang
HIGHLIGHT: In this work, we investigate the biases present in Hindi language representations such as caste and religion associated biases.

- 77, TITLE: AmbiPun: Generating Humorous Puns with Ambiguous Context
<https://aclanthology.org/2022.naacl-main.77>
AUTHORS: Anirudh Mittal, Yufei Tian, Nanyun Peng
HIGHLIGHT: In this paper, we propose a simple yet effective way to generate pun sentences that does not require any training on existing puns.
- 78, TITLE: EmpHi: Generating Empathetic Responses with Human-like Intents
<https://aclanthology.org/2022.naacl-main.78>
AUTHORS: Mao Yan Chen, Siheng Li, Yujiu Yang
HIGHLIGHT: To address the bias of the empathetic intents distribution between empathetic dialogue models and humans, we propose a novel model to generate empathetic responses with human-consistent empathetic intents, EmpHi for short.
- 79, TITLE: Yes, No or IDK: The Challenge of Unanswerable Yes/No Questions
<https://aclanthology.org/2022.naacl-main.79>
AUTHORS: Elior Sulem, Jamaal Hay, Dan Roth
HIGHLIGHT: In this paper, we extend the Yes/No QA task, adding questions with an IDK answer, and show its considerable difficulty compared to the original 2-label task.
- 80, TITLE: Inducing and Using Alignments for Transition-based AMR Parsing
<https://aclanthology.org/2022.naacl-main.80>
AUTHORS: Andrew Drozdov, Jiawei Zhou, Radu Florian, Andrew McCallum, Tahira Naseem, Yoon Kim, Ram?n Astudillo
HIGHLIGHT: Parsers also train on a point-estimate of the alignment pipeline, neglecting the uncertainty due to the inherent ambiguity of alignment. In this work we explore two avenues for overcoming these limitations.
- 81, TITLE: Masked Part-Of-Speech Model: Does Modeling Long Context Help Unsupervised POS-tagging?
<https://aclanthology.org/2022.naacl-main.81>
AUTHORS: Xiang Zhou, Shiyue Zhang, Mohit Bansal
HIGHLIGHT: To facilitate flexible dependency modeling, we propose a Masked Part-of-Speech Model (MPoSM), inspired by the recent success of Masked Language Models (MLM).
- 82, TITLE: DREAM: Improving Situational QA by First Elaborating the Situation
<https://aclanthology.org/2022.naacl-main.82>
AUTHORS: Yuling Gu, Bhavana Dalvi, Peter Clark
HIGHLIGHT: While we do not know how language models (LMs) answer such questions, we conjecture that they may answer more accurately if they are also provided with additional details about the question situation, elaborating the "scene". To test this conjecture, we train a new model, DREAM, to answer questions that elaborate the scenes that situated questions are about, and then provide those elaborations as additional context to a question-answering (QA) model.
- 83, TITLE: CoSe-Co: Text Conditioned Generative Commonsense Contextualizer
<https://aclanthology.org/2022.naacl-main.83>
AUTHORS: Rachit Bansal, Milan Aggarwal, Sumit Bhatia, Jivat Kaur, Balaji Krishnamurthy
HIGHLIGHT: However, training on symbolic KG entities limits their applicability in tasks involving natural language text where they ignore overall context. To mitigate this, we propose a Commonsense Contextualizer (CoSe-Co) conditioned on sentences as input to make it generically usable in tasks for generating knowledge relevant to the overall context of input text.
- 84, TITLE: Probing via Prompting
<https://aclanthology.org/2022.naacl-main.84>
AUTHORS: Jiaoda Li, Ryan Cotterell, Mrinmaya Sachan
HIGHLIGHT: However, the mechanism of selecting the probe model has recently been subject to intense debate, as it is not clear if the probes are merely extracting information or modelling the linguistic property themselves. To address this challenge, this paper introduces a novel model-free approach to probing via prompting, which formulates probing as a prompting task.
- 85, TITLE: Database Search Results Disambiguation for Task-Oriented Dialog Systems
<https://aclanthology.org/2022.naacl-main.85>
AUTHORS: Kun Qian, Satwik Kottur, Ahmad Beirami, Shahin Shayandeh, Paul Crook, Alborz Geramifard, Zhou Yu, Chinnadhurai Sankar
HIGHLIGHT: In this paper, we propose Database Search Result (DSR) Disambiguation, a novel task that focuses on disambiguating database search results, which enhances user experience by allowing them to choose from multiple options instead of just one.
- 86, TITLE: Unsupervised Slot Schema Induction for Task-oriented Dialog
<https://aclanthology.org/2022.naacl-main.86>
AUTHORS: Dian Yu, Mingqiu Wang, Yuan Cao, Izhak Shafran, Laurent Shafey, Hagen Soltau
HIGHLIGHT: In practical applications, manually designing schemas can be error-prone, laborious, iterative, and slow, especially when the schema is complicated. To alleviate this expensive and time consuming process, we propose an unsupervised approach for slot schema induction from unlabeled dialog corpora.
- 87, TITLE: Towards a Progression-Aware Autonomous Dialogue Agent
<https://aclanthology.org/2022.naacl-main.87>

AUTHORS: Abraham Sanders, Tomek Strzalkowski, Mei Si, Albert Chang, Deepanshu Dey, Jonas Braasch, Dakuo Wang
HIGHLIGHT: Thus, we propose a framework in which dialogue agents can evaluate the progression of a conversation toward or away from desired outcomes, and use this signal to inform planning for subsequent responses.

88, TITLE: Cross-Domain Detection of GPT-2-Generated Technical Text
<https://aclanthology.org/2022.naacl-main.88>
AUTHORS: Juan Rodriguez, Todd Hay, David Gros, Zain Shamsi, Ravi Srinivasan
HIGHLIGHT: In this paper we examine the problem of detecting GPT-2-generated technical research text.

89, TITLE: DISAPERE: A Dataset for Discourse Structure in Peer Review Discussions
<https://aclanthology.org/2022.naacl-main.89>
AUTHORS: Neha Kennard, Tim O’Gorman, Rajarshi Das, Akshay Sharma, Chhandak Bagchi, Matthew Clinton, Pranay Kumar Yelugam, Hamed Zamani, Andrew McCallum
HIGHLIGHT: We present DISAPERE, a labeled dataset of 20k sentences contained in 506 review-rebuttal pairs in English, annotated by experts.

90, TITLE: MultiSpanQA: A Dataset for Multi-Span Question Answering
<https://aclanthology.org/2022.naacl-main.90>
AUTHORS: Haonan Li, Martin Tomko, Maria Vasardani, Timothy Baldwin
HIGHLIGHT: In this paper, we present MultiSpanQA, a new dataset that focuses on multi-span questions.

91, TITLE: Context-Aware Abbreviation Expansion Using Large Language Models
<https://aclanthology.org/2022.naacl-main.91>
AUTHORS: Shanqing Cai, Subhashini Venugopalan, Katrin Tomanek, Ajit Narayanan, Meredith Morris, Michael Brenner
HIGHLIGHT: Motivated by the need for accelerating text entry in augmentative and alternative communication (AAC) for people with severe motor impairments, we propose a paradigm in which phrases are abbreviated aggressively as primarily word-initial letters.

92, TITLE: Theory-Grounded Measurement of U.S. Social Stereotypes in English Language Models
<https://aclanthology.org/2022.naacl-main.92>
AUTHORS: Yang Cao, Anna Sotnikova, Hal Daum? III, Rachel Rudinger, Linda Zou
HIGHLIGHT: We introduce the sensitivity test (SeT) for measuring stereotypical associations from language models.

93, TITLE: Sort by Structure: Language Model Ranking as Dependency Probing
<https://aclanthology.org/2022.naacl-main.93>
AUTHORS: Max Miller-Eberstein, Rob Goot, Barbara Plank
HIGHLIGHT: We propose probing to rank LMs, specifically for parsing dependencies in a given language, by measuring the degree to which labeled trees are recoverable from an LM’s contextualized embeddings.

94, TITLE: Quantifying Synthesis and Fusion and their Impact on Machine Translation
<https://aclanthology.org/2022.naacl-main.94>
AUTHORS: Arturo Oncevay, Duygu Ataman, Niels Van Berkel, Barry Haddow, Alexandra Birch, Johannes Bjerva
HIGHLIGHT: However, literature in Natural Language Processing (NLP) typically labels a whole language with a strict type of morphology, e.g. fusional or agglutinative. In this work, we propose to reduce the rigidity of such claims, by quantifying morphological typology at the word and segment level.

95, TITLE: Commonsense and Named Entity Aware Knowledge Grounded Dialogue Generation
<https://aclanthology.org/2022.naacl-main.95>
AUTHORS: Deeksha Varshney, Akshara Prabhakar, Asif Ekbal
HIGHLIGHT: In this paper, we present a novel open-domain dialogue generation model which effectively utilizes the large-scale commonsense and named entity based knowledge in addition to the unstructured topic-specific knowledge associated with each utterance.

96, TITLE: Efficient Hierarchical Domain Adaptation for Pretrained Language Models
<https://aclanthology.org/2022.naacl-main.96>
AUTHORS: Alexandra Chronopoulou, Matthew Peters, Jesse Dodge
HIGHLIGHT: In this paper, we introduce a method to permit domain adaptation to many diverse domains using a computationally efficient adapter approach.

97, TITLE: Hatemoji: A Test Suite and Adversarially-Generated Dataset for Benchmarking and Detecting Emoji-Based Hate
<https://aclanthology.org/2022.naacl-main.97>
AUTHORS: Hannah Kirk, Bertie Vidgen, Paul Rottger, Tristan Thrush, Scott Hale
HIGHLIGHT: We present HatemojiCheck, a test suite of 3,930 short-form statements that allows us to evaluate performance on hateful language expressed with emoji.

98, TITLE: On the Economics of Multilingual Few-shot Learning: Modeling the Cost-Performance Trade-offs of Machine Translated and Manual Data
<https://aclanthology.org/2022.naacl-main.98>
AUTHORS: Kabir Ahuja, Monojit Choudhury, Sandipan Dandapat

HIGHLIGHT: Borrowing ideas from Production functions in micro-economics, in this paper we introduce a framework to systematically evaluate the performance and cost trade-offs between machine-translated and manually-created labelled data for task-specific fine-tuning of massively multilingual language models.

99, **TITLE:** Learning to Selectively Learn for Weakly Supervised Paraphrase Generation with Model-based Reinforcement Learning

<https://aclanthology.org/2022.naacl-main.99>

AUTHORS: Haiyan Yin, Dingcheng Li, Ping Li

HIGHLIGHT: While data selection is privileged for the target task which has noisy data, developing a reinforced selective learning regime faces several unresolved challenges. In this paper, we carry on important discussions about the above problem and present a new model that could partially overcome the discussed issues with a model-based planning feature and a reward normalization feature.

100, **TITLE:** Quality-Aware Decoding for Neural Machine Translation

<https://aclanthology.org/2022.naacl-main.100>

AUTHORS: Patrick Fernandes, Ant?nio Farinhas, Ricardo Rei, Jos? De Souza, Perez Ogayo, Graham Neubig, Andre Martins

HIGHLIGHT: Despite the progress in machine translation quality estimation and evaluation in the last years, decoding in neural machine translation (NMT) is mostly oblivious to this and centers around finding the most probable translation according to the model (MAP decoding), approximated with beam search. In this paper, we bring together these two lines of research and propose quality-aware decoding for NMT, by leveraging recent breakthroughs in reference-free and reference-based MT evaluation through various inference methods like N-best reranking and minimum Bayes risk decoding.

101, **TITLE:** Pretrained Models for Multilingual Federated Learning

<https://aclanthology.org/2022.naacl-main.101>

AUTHORS: Orion Weller, Marc Marone, Vladimir Braverman, Dawn Lawrie, Benjamin Van Durme

HIGHLIGHT: We explore three multilingual language tasks, language modeling, machine translation, and text classification using differing federated and non-federated learning algorithms.

102, **TITLE:** AcTune: Uncertainty-Based Active Self-Training for Active Fine-Tuning of Pretrained Language Models

<https://aclanthology.org/2022.naacl-main.102>

AUTHORS: Yue Yu, Ling kai Kong, Jieyu Zhang, Rongzhi Zhang, Chao Zhang

HIGHLIGHT: We develop AcTune, a new framework that improves the label efficiency of active PLM fine-tuning by unleashing the power of unlabeled data via self-training.

103, **TITLE:** Label Anchored Contrastive Learning for Language Understanding

<https://aclanthology.org/2022.naacl-main.103>

AUTHORS: Zhenyu Zhang, Yuming Zhao, Meng Chen, Xiaodong He

HIGHLIGHT: Intuitively, the class label itself has the intrinsic ability to perform hard positive/negative mining, which is crucial for CL. Motivated by this, we propose a novel label anchored contrastive learning approach (denoted as LaCon) for language understanding.

104, **TITLE:** Go Back in Time: Generating Flashbacks in Stories with Event Temporal Prompts

<https://aclanthology.org/2022.naacl-main.104>

AUTHORS: Rujun Han, Hong Chen, Yufei Tian, Nanyun Peng

HIGHLIGHT: Two major issues in existing systems that exacerbate the challenges: 1) temporal bias in pertaining and story datasets that leads to monotonic event temporal orders; 2) lack of explicit guidance that helps machines decide where to insert *flashbacks*. We propose to address these issues using structured storylines to encode events and their pair-wise temporal relations (before, after and vague) as **temporal prompts** that guide how stories should unfold temporally.

105, **TITLE:** Forecasting COVID-19 Caseloads Using Unsupervised Embedding Clusters of Social Media Posts

<https://aclanthology.org/2022.naacl-main.105>

AUTHORS: Felix Drinkall, Stefan Zohren, Janet Pierrehumbert

HIGHLIGHT: We present a novel approach incorporating transformer-based language models into infectious disease modelling.

106, **TITLE:** Many Hands Make Light Work: Using Essay Traits to Automatically Score Essays

<https://aclanthology.org/2022.naacl-main.106>

AUTHORS: Rahul Kumar, Sandeep Mathias, Sriparna Saha, Pushpak Bhattacharyya

HIGHLIGHT: In this paper, we describe a way to score essays using a multi-task learning (MTL) approach, where scoring the essay holistically is the primary task, and scoring the essay traits is the auxiliary task.

107, **TITLE:** Natural Language Inference with Self-Attention for Veracity Assessment of Pandemic Claims

<https://aclanthology.org/2022.naacl-main.107>

AUTHORS: Miguel Arana-Catania, Elena Kochkina, Arkaitz Zubiaga, Maria Liakata, Robert Procter, Yulan He

HIGHLIGHT: We present a comprehensive work on automated veracity assessment from dataset creation to developing novel methods based on Natural Language Inference (NLI), focusing on misinformation related to the COVID-19 pandemic.

108, **TITLE:** Beyond Emotion: A Multi-Modal Dataset for Human Desire Understanding

<https://aclanthology.org/2022.naacl-main.108>

- AUTHORS:** Ao Jia, Yu He, Yazhou Zhang, Sagar Uprety, Dawei Song, Christina Lioma
HIGHLIGHT: As a strikingly understudied task, it is difficult for machines to model and understand desire due to the unavailability of benchmarking datasets with desire and emotion labels. To bridge this gap, we present MSED, the first multi-modal and multi-task sentiment, emotion and desire dataset, which contains 9,190 text-image pairs, with English text.
- 109, **TITLE:** Relation-Specific Attentions over Entity Mentions for Enhanced Document-Level Relation Extraction
<https://aclanthology.org/2022.naacl-main.109>
AUTHORS: Jiaxin Yu, Deqing Yang, Shuyu Tian
HIGHLIGHT: As a result, the distinct semantics between the different mentions of an entity are overlooked. To address this problem, we propose RSMAN in this paper which performs selective attentions over different entity mentions with respect to candidate relations.
- 110, **TITLE:** Twitter-COMMs: Detecting Climate, COVID, and Military Multimodal Misinformation
<https://aclanthology.org/2022.naacl-main.110>
AUTHORS: Giscard Biamby, Grace Luo, Trevor Darrell, Anna Rohrbach
HIGHLIGHT: Detecting out-of-context media, such as "misp captioned" images on Twitter, is a relevant problem, especially in domains of high public significance. In this work we aim to develop defenses against such misinformation for the topics of Climate Change, COVID-19, and Military Vehicles.
- 111, **TITLE:** BlonDe: An Automatic Evaluation Metric for Document-level Machine Translation
<https://aclanthology.org/2022.naacl-main.111>
AUTHORS: Yuchen Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan, Ming Zhou
HIGHLIGHT: This paper introduces a novel automatic metric BlonDe to widen the scope of automatic MT evaluation from sentence to document level.
- 112, **TITLE:** Disentangled Learning of Stance and Aspect Topics for Vaccine Attitude Detection in Social Media
<https://aclanthology.org/2022.naacl-main.112>
AUTHORS: Lixing Zhu, Zheng Fang, Gabriele Pergola, Robert Procter, Yulan He
HIGHLIGHT: Instead, with the aim of leveraging the large amount of unannotated data now available on vaccination, we propose a novel semi-supervised approach for vaccine attitude detection, called VADet.
- 113, **TITLE:** SKILL: Structured Knowledge Infusion for Large Language Models
<https://aclanthology.org/2022.naacl-main.113>
AUTHORS: Fedor Moiseev, Zhe Dong, Enrique Alfonseca, Martin Jaggi
HIGHLIGHT: In this work, we propose a method to infuse structured knowledge into LLMs, by directly training T5 models on factual triples of knowledge graphs (KGs).
- 114, **TITLE:** Same Neurons, Different Languages: Probing Morphosyntax in Multilingual Pre-trained Models
<https://aclanthology.org/2022.naacl-main.114>
AUTHORS: Karolina Stanczak, Edoardo Ponti, Lucas Torroba Hennigen, Ryan Cotterell, Isabelle Augenstein
HIGHLIGHT: In this work, we conjecture that multilingual pre-trained models can derive language-universal abstractions about grammar.
- 115, **TITLE:** Aspect Is Not You Need: No-aspect Differential Sentiment Framework for Aspect-based Sentiment Analysis
<https://aclanthology.org/2022.naacl-main.115>
AUTHORS: Jiahao Cao, Rui Liu, Huailiang Peng, Lei Jiang, Xu Bai
HIGHLIGHT: In this work, we analyze the ABSA task from a novel cognition perspective: humans can often judge the sentiment of an aspect even if they do not know what the aspect is.
- 116, **TITLE:** MoEBERT: from BERT to Mixture-of-Experts via Importance-Guided Adaptation
<https://aclanthology.org/2022.naacl-main.116>
AUTHORS: Simiao Zuo, Qingru Zhang, Chen Liang, Pengcheng He, Tuo Zhao, Weizhu Chen
HIGHLIGHT: We propose MoEBERT, which uses a Mixture-of-Experts structure to increase model capacity and inference speed.
- 117, **TITLE:** Implicit n-grams Induced by Recurrence
<https://aclanthology.org/2022.naacl-main.117>
AUTHORS: Xiaobing Sun, Wei Lu
HIGHLIGHT: In this work, we present a study that shows there actually exist some explainable components that reside within the hidden states, which are reminiscent of the classical n-grams features.
- 118, **TITLE:** Guiding Visual Question Generation
<https://aclanthology.org/2022.naacl-main.118>
AUTHORS: Nihir Vedh, Zixu Wang, Marek Rei, Yishu Miao, Lucia Specia
HIGHLIGHT: We present Guiding Visual Question Generation - a variant of VQG which conditions the question generator on categorical information based on expectations on the type of question and the objects it should explore.
- 119, **TITLE:** OPERA: Operation-Pivoted Discrete Reasoning over Text
<https://aclanthology.org/2022.naacl-main.119>

- AUTHORS:** Yongwei Zhou, Junwei Bao, Chaoqun Duan, Haipeng Sun, Jiahui Liang, Yifan Wang, Jing Zhao, Youzheng Wu, Xiaodong He, Tiejun Zhao
HIGHLIGHT: However, they ignore the utilization of symbolic operations and encounter a lack of reasoning ability and interpretability. To inherit the advantages of these two types of methods, we propose OPERA, an operation-pivoted discrete reasoning framework, where lightweight symbolic operations (compared with logical forms) as neural modules are utilized to facilitate the reasoning ability and interpretability.
- 120, **TITLE:** Improving Multi-Document Summarization through Referenced Flexible Extraction with Credit-Awareness
<https://aclanthology.org/2022.naacl-main.120>
AUTHORS: Yun-Zhu Song, Yi-Syuan Chen, Hong-Han Shuai
HIGHLIGHT: In this paper, we present an extract-then-abstract Transformer framework to overcome the problem.
- 121, **TITLE:** Improving Constituent Representation with Hypertree Neural Networks
<https://aclanthology.org/2022.naacl-main.121>
AUTHORS: Hao Zhou, Gongshen Liu, Kewei Tu
HIGHLIGHT: In this paper, we aim to improve representations of constituent spans using a novel hypertree neural networks (HTNN) that is structured with constituency parse trees.
- 122, **TITLE:** Measuring Fairness with Biased Rulers: A Comparative Study on Bias Metrics for Pre-trained Language Models
<https://aclanthology.org/2022.naacl-main.122>
AUTHORS: Pieter Delobelle, Ewoenam Tokpo, Toon Calders, Bettina Berendt
HIGHLIGHT: We survey the literature on fairness metrics for pre-trained language models and experimentally evaluate compatibility, including both biases in language models and in their downstream tasks.
- 123, **TITLE:** MuCPAD: A Multi-Domain Chinese Predicate-Argument Dataset
<https://aclanthology.org/2022.naacl-main.123>
AUTHORS: Yahui Liu, Haoping Yang, Chen Gong, Qingrong Xia, Zhenghua Li, Min Zhang
HIGHLIGHT: In order to facilitate research on cross-domain SRL, this paper presents MuCPAD, a multi-domain Chinese predicate-argument dataset, which consists of 30,897 sentences and 92,051 predicates from six different domains.
- 124, **TITLE:** Representation Learning for Conversational Data using Discourse Mutual Information Maximization
<https://aclanthology.org/2022.naacl-main.124>
AUTHORS: Bishal Santra, Sumegh Roychowdhury, Aishik Mandal, Vasu Gurram, Atharva Naik, Manish Gupta, Pawan Goyal
HIGHLIGHT: Hence, we propose a structure-aware Mutual Information based loss-function DMI (Discourse Mutual Information) for training dialog-representation models, that additionally captures the inherent uncertainty in response prediction.
- 125, **TITLE:** ValCAT: Variable-Length Contextualized Adversarial Transformations Using Encoder-Decoder Language Model
<https://aclanthology.org/2022.naacl-main.125>
AUTHORS: Chuyun Deng, Mingxuan Liu, Yue Qin, Jia Zhang, Hai-Xin Duan, Donghong Sun
HIGHLIGHT: In this paper, we propose ValCAT, a black-box attack framework that misleads the language model by applying variable-length contextualized transformations to the original text.
- 126, **TITLE:** A Study of Syntactic Multi-Modality in Non-Autoregressive Machine Translation
<https://aclanthology.org/2022.naacl-main.126>
AUTHORS: Kexun Zhang, Rui Wang, Xu Tan, Junliang Guo, Yi Ren, Tao Qin, Tie-Yan Liu
HIGHLIGHT: In this paper, we conduct a systematic study on the syntactic multi-modality problem.
- 127, **TITLE:** CIAug: Equipping Interpolative Augmentation with Curriculum Learning
<https://aclanthology.org/2022.naacl-main.127>
AUTHORS: Ramit Sawhney, Ritesh Soun, Shrey Pandit, Megh Thakkar, Sarvagya Malaviya, Yuval Pinter
HIGHLIGHT: We propose CIAug, a novel curriculum-based learning method that builds upon mixup.
- 128, **TITLE:** Proposition-Level Clustering for Multi-Document Summarization
<https://aclanthology.org/2022.naacl-main.128>
AUTHORS: Ori Ernst, Avi Caciularu, Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Jacob Goldberger, Ido Dagan
HIGHLIGHT: In this work, we revisit the clustering approach, grouping together sub-sentential propositions, aiming at more precise information alignment.
- 129, **TITLE:** Non-Autoregressive Machine Translation: It's Not as Fast as it Seems
<https://aclanthology.org/2022.naacl-main.129>
AUTHORS: Jindrich Helcl, Barry Haddow, Alexandra Birch
HIGHLIGHT: In this paper, we point out flaws in the evaluation methodology present in the literature on NAR models and we provide a fair comparison between a state-of-the-art NAR model and the autoregressive submissions to the shared task.
- 130, **TITLE:** BAD-X: Bilingual Adapters Improve Zero-Shot Cross-Lingual Transfer
<https://aclanthology.org/2022.naacl-main.130>
AUTHORS: Marinela Parovic, Goran Glava?, Ivan Vulic, Anna Korhonen

HIGHLIGHT: In this work, we show that it is more effective to learn bilingual language pair adapters (BAs) when the goal is to optimize performance for a particular source-target transfer direction.

131, **TITLE:** Combining Humor and Sarcasm for Improving Political Parody Detection

<https://aclanthology.org/2022.naacl-main.131>

AUTHORS: Xiao Ao, Danae Sanchez Villegas, Daniel Preotiuc-Pietro, Nikolaos Aletras

HIGHLIGHT: This paper explores jointly modelling these figurative tropes with the goal of improving performance of political parody detection in tweets. To this end, we present a multi-encoder model that combines three parallel encoders to enrich parody-specific representations with humor and sarcasm information.

132, **TITLE:** TIE: Topological Information Enhanced Structural Reading Comprehension on Web Pages

<https://aclanthology.org/2022.naacl-main.132>

AUTHORS: Zihan Zhao, Lu Chen, Ruisheng Cao, Hongshen Xu, Xingyu Chen, Kai Yu

HIGHLIGHT: In this work, we propose a Topological Information Enhanced model (TIE), which transforms the token-level task into a tag-level task by introducing a two-stage process (i.e. node locating and answer refining).

133, **TITLE:** RSTGen: Imbuing Fine-Grained Interpretable Control into Long-FormText Generators

<https://aclanthology.org/2022.naacl-main.133>

AUTHORS: Rilwan Adewoyin, Ritabrata Dutta, Yulan He

HIGHLIGHT: In this paper, we study the task of improving the cohesion and coherence of long-form text generated by language models.

134, **TITLE:** Intent Detection and Discovery from User Logs via Deep Semi-Supervised Contrastive Clustering

<https://aclanthology.org/2022.naacl-main.134>

AUTHORS: Rajat Kumar, Mayur Patidar, Vaibhav Varshney, Lovekesh Vig, Gautam Shroff

HIGHLIGHT: Unlike existing approaches that rely on epoch wise cluster alignment, we propose an end-to-end deep contrastive clustering algorithm that jointly updates model parameters and cluster centers via supervised and self-supervised learning and optimally utilizes both labeled and unlabeled data.

135, **TITLE:** Extending Multi-Text Sentence Fusion Resources via Pyramid Annotations

<https://aclanthology.org/2022.naacl-main.135>

AUTHORS: Daniela Brook Weiss, Paul Roit, Ori Ernst, Ido Dagan

HIGHLIGHT: In this paper, we revisit and substantially extend previous dataset creation efforts.

136, **TITLE:** The Devil is in the Details: On the Pitfalls of Vocabulary Selection in Neural Machine Translation

<https://aclanthology.org/2022.naacl-main.136>

AUTHORS: Tobias Domhan, Eva Hasler, Ke Tran, Sony Trenous, Bill Byrne, Felix Hieber

HIGHLIGHT: We propose a model of vocabulary selection, integrated into the neural translation model, that predicts the set of allowed output words from contextualized encoder representations.

137, **TITLE:** MultiCite: Modeling realistic citations requires moving beyond the single-sentence single-label setting

<https://aclanthology.org/2022.naacl-main.137>

AUTHORS: Anne Lauscher, Brandon Ko, Bailey Kuehl, Sophie Johnson, Arman Cohan, David Jurgens, Kyle Lo

HIGHLIGHT: Yet, recent work in CCA is often approached as a single-sentence, single-label classification task, and thus many datasets used to develop modern computational approaches fail to capture this interesting discourse. To address this research gap, we highlight three understudied phenomena for CCA and release MULTICITE, a new dataset of 12.6K citation contexts from 1.2K computational linguistics papers that fully models these phenomena.

138, **TITLE:** DEGREE: A Data-Efficient Generation-Based Event Extraction Model

<https://aclanthology.org/2022.naacl-main.138>

AUTHORS: I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, Nanyun Peng

HIGHLIGHT: In this paper, we focus on low-resource end-to-end event extraction and propose DEGREE, a data-efficient model that formulates event extraction as a conditional generation problem.

139, **TITLE:** Bridging the Gap between Language Models and Cross-Lingual Sequence Labeling

<https://aclanthology.org/2022.naacl-main.139>

AUTHORS: Nuo Chen, Linjun Shou, Ming Gong, Jian Pei, Daxin Jiang

HIGHLIGHT: In this paper, we first design a pre-training task tailored for xSL named Cross-lingual Language Informative Span Masking (CLISM) to eliminate the objective gap in a self-supervised manner. Second, we present ContrAstive-Consistency Regularization (CACR), which utilizes contrastive learning to encourage the consistency between representations of input parallel sequences via unsupervised cross-lingual instance-wise training signals during pre-training.

140, **TITLE:** Hero-Gang Neural Model For Named Entity Recognition

<https://aclanthology.org/2022.naacl-main.140>

AUTHORS: Jinpeng Hu, Yaling Shen, Yang Liu, Xiang Wan, Tsung-Hui Chang

HIGHLIGHT: Unfortunately, although these models can capture effective global context information, they are still limited in the local feature and position information extraction, which is critical in NER. In this paper, to address this limitation, we propose a novel Hero-Gang Neural structure (HGN), including the Hero and Gang module, to leverage both global and local information to promote NER.

- 141, TITLE: MGIMN: Multi-Grained Interactive Matching Network for Few-shot Text Classification
<https://aclanthology.org/2022.naacl-main.141>
AUTHORS: Jianhai Zhang, Mieradilijiang Maimaiti, Gao Xing, Yuanhang Zheng, Ji Zhang
HIGHLIGHT: To deal with these issues, we propose a meta-learning based method MGIMN which performs instance-wise comparison followed by aggregation to generate class-wise matching vectors instead of prototype learning.
- 142, TITLE: All You May Need for VQA are Image Captions
<https://aclanthology.org/2022.naacl-main.142>
AUTHORS: Soravit Changpinyo, Doron Kukliansy, Idan Szpektor, Xi Chen, Nan Ding, Radu Soricut
HIGHLIGHT: In this paper, we propose a method that automatically derives VQA examples at volume, by leveraging the abundance of existing image-caption annotations combined with neural models for textual question generation.
- 143, TITLE: Frustratingly Easy System Combination for Grammatical Error Correction
<https://aclanthology.org/2022.naacl-main.143>
AUTHORS: Muhammad Qorib, Seung-Hoon Na, Hwee Tou Ng
HIGHLIGHT: In this paper, we formulate system combination for grammatical error correction (GEC) as a simple machine learning task: binary classification.
- 144, TITLE: Simple Local Attentions Remain Competitive for Long-Context Tasks
<https://aclanthology.org/2022.naacl-main.144>
AUTHORS: Wenhan Xiong, Barlas Oguz, Anchit Gupta, Xilun Chen, Diana Liskovich, Omer Levy, Scott Yih, Yashar Mehdad
HIGHLIGHT: Despite the abundance of research along this direction, it is still difficult to gauge the relative effectiveness of these models in practical use cases, e.g., if we apply these models following the pretrain-and-finetune paradigm. In this work, we aim to conduct a thorough analysis of these emerging models with large-scale and controlled experiments.
- 145, TITLE: Even the Simplest Baseline Needs Careful Re-investigation: A Case Study on XML-CNN
<https://aclanthology.org/2022.naacl-main.145>
AUTHORS: Si-An Chen, Jie-jyun Liu, Tsung-Han Yang, Hsuan-Tien Lin, Chih-Jen Lin
HIGHLIGHT: In this work, through an astonishing example we argue that more efforts should be paid to ensure the progress in developing a new deep learning method.
- 146, TITLE: Multi-Relational Graph Transformer for Automatic Short Answer Grading
<https://aclanthology.org/2022.naacl-main.146>
AUTHORS: Rajat Agarwal, Varun Khurana, Karish Grover, Mukesh Mohania, Vikram Goyal
HIGHLIGHT: Most existing methods utilize sequential context to compare two sentences and ignore the structural context of the sentence; therefore, these methods may not result in the desired performance. In this paper, we overcome this problem by proposing a Multi-Relational Graph Transformer, MitiGaTe, to prepare token representations considering the structural context.
- 147, TITLE: Event Schema Induction with Double Graph Autoencoders
<https://aclanthology.org/2022.naacl-main.147>
AUTHORS: Xiaomeng Jin, Manling Li, Heng Ji
HIGHLIGHT: We propose a new event schema induction framework using double graph autoencoders, which captures the global dependencies among nodes in event graphs.
- 148, TITLE: CS1QA: A Dataset for Assisting Code-based Question Answering in an Introductory Programming Course
<https://aclanthology.org/2022.naacl-main.148>
AUTHORS: Changyoon Lee, Yeon Seonwoo, Alice Oh
HIGHLIGHT: We introduce CS1QA, a dataset for code-based question answering in the programming education domain.
- 149, TITLE: Unsupervised Cross-Lingual Transfer of Structured Predictors without Source Data
<https://aclanthology.org/2022.naacl-main.149>
AUTHORS: Kemal Kurniawan, Lea Freemann, Philip Schulz, Trevor Cohn
HIGHLIGHT: To that end, we generalise methods for unsupervised transfer from multiple input models for structured prediction.
- 150, TITLE: Don't Take It Literally: An Edit-Invariant Sequence Loss for Text Generation
<https://aclanthology.org/2022.naacl-main.150>
AUTHORS: Guangyi Liu, Zichao Yang, Tianhua Tao, Xiaodan Liang, Junwei Bao, Zhen Li, Xiaodong He, Shuguang Cui, Zhiting Hu
HIGHLIGHT: To address the challenge, we propose a novel Edit-Invariant Sequence Loss (EISL), which computes the matching loss of a target n-gram with all n-grams in the generated sequence.
- 151, TITLE: Modeling Exemplification in Long-form Question Answering via Retrieval
<https://aclanthology.org/2022.naacl-main.151>
AUTHORS: Shufan Wang, Fangyuan Xu, Laure Thompson, Eunsol Choi, Mohit Iyyer
HIGHLIGHT: In this paper, we provide the first computational study of exemplification in QA, performing a fine-grained annotation of different types of examples (e.g., hypotheticals, anecdotes) in three corpora.

- 152, TITLE: D2U: Distance-to-Uniform Learning for Out-of-Scope Detection
<https://aclanthology.org/2022.naacl-main.152>
AUTHORS: Eyup Yilmaz, Cagri Toraman
HIGHLIGHT: Specifically, we propose a zero-shot post-processing step, called Distance-to-Uniform (D2U), exploiting not only the classification confidence score, but the shape of the entire output distribution.
- 153, TITLE: Reference-free Summarization Evaluation via Semantic Correlation and Compression Ratio
<https://aclanthology.org/2022.naacl-main.153>
AUTHORS: Yizhu Liu, Qi Jia, Kenny Zhu
HIGHLIGHT: In this paper, we propose a new automatic reference-free evaluation metric that compares semantic distribution between source document and summary by pretrained language models and considers summary compression ratio.
- 154, TITLE: KroneckerBERT: Significant Compression of Pre-trained Language Models Through Kronecker Decomposition and Knowledge Distillation
<https://aclanthology.org/2022.naacl-main.154>
AUTHORS: Marzieh Tahaei, Ella Charlaix, Vahid Nia, Ali Ghodsi, Mehdi Rezagholizadeh
HIGHLIGHT: We present our KroneckerBERT, a compressed version of the BERT_BASE model obtained by compressing the embedding layer and the linear mappings in the multi-head attention, and the feed-forward network modules in the Transformer layers.
- 155, TITLE: Building a Role Specified Open-Domain Dialogue System Leveraging Large-Scale Language Models
<https://aclanthology.org/2022.naacl-main.155>
AUTHORS: Sanghwan Bae, Donghyun Kwak, Sungdong Kim, Donghoon Ham, Soyoung Kang, Sang-Woo Lee, Woomyoung Park
HIGHLIGHT: In this work, we study the challenge of imposing roles on open-domain dialogue systems, with the goal of making the systems maintain consistent roles while conversing naturally with humans.
- 156, TITLE: Sentence-Level Resampling for Named Entity Recognition
<https://aclanthology.org/2022.naacl-main.156>
AUTHORS: Xiaochen Wang, Yue Wang
HIGHLIGHT: To alleviate data imbalance, we propose a set of sentence-level resampling methods where the importance of each training sentence is computed based on its tokens and entities.
- 157, TITLE: Word Tour: One-dimensional Word Embeddings via the Traveling Salesman Problem
<https://aclanthology.org/2022.naacl-main.157>
AUTHORS: Ryoma Sato
HIGHLIGHT: In this study, we propose WordTour, unsupervised one-dimensional word embeddings.
- 158, TITLE: On the Diversity and Limits of Human Explanations
<https://aclanthology.org/2022.naacl-main.158>
AUTHORS: Chenhao Tan
HIGHLIGHT: Our goal is to provide an overview of the diversity of explanations, discuss human limitations in providing explanations, and ultimately provide implications for collecting and using human explanations in NLP.
- 159, TITLE: Locally Aggregated Feature Attribution on Natural Language Model Understanding
<https://aclanthology.org/2022.naacl-main.159>
AUTHORS: Sheng Zhang, Jin Wang, Haitao Jiang, Rui Song
HIGHLIGHT: In this work, we propose Locally Aggregated Feature Attribution (LAFA), a novel gradient-based feature attribution method for NLP models.
- 160, TITLE: Generic and Trend-aware Curriculum Learning for Relation Extraction
<https://aclanthology.org/2022.naacl-main.160>
AUTHORS: Nidhi Vakil, Hadi Amiri
HIGHLIGHT: We present a generic and trend-aware curriculum learning approach that effectively integrates textual and structural information in text graphs for relation extraction between entities, which we consider as node pairs in graphs.
- 161, TITLE: On Systematic Style Differences between Unsupervised and Supervised MT and an Application for High-Resource Machine Translation
<https://aclanthology.org/2022.naacl-main.161>
AUTHORS: Kelly Marchisio, Markus Freitag, David Grangier
HIGHLIGHT: We compare translations from supervised and unsupervised MT systems of similar quality, finding that unsupervised output is more fluent and more structurally different in comparison to human translation than is supervised MT. We then demonstrate a way to combine the benefits of both methods into a single system which results in improved adequacy and fluency as rated by human evaluators.
- 162, TITLE: Evidentiality-guided Generation for Knowledge-Intensive NLP Tasks
<https://aclanthology.org/2022.naacl-main.162>
AUTHORS: Akari Asai, Matt Gardner, Hannaneh Hajishirzi
HIGHLIGHT: This work introduces a method to incorporate evidentiality of passages-whether a passage contains correct evidence to support the output-into training the generator.

- 163, TITLE: Modularized Transfer Learning with Multiple Knowledge Graphs for Zero-shot Commonsense Reasoning
<https://aclanthology.org/2022.naacl-main.163>
AUTHORS: Yu Jin Kim, Beong-woo Kwak, Youngwook Kim, Reinald Kim Amplayo, Seung-won Hwang, Jinyoung Yeo
HIGHLIGHT: Considering the increasing type of different commonsense KGs, this paper aims to extend the zero-shot transfer learning scenario into multiple-source settings, where different KGs can be utilized synergetically.
- 164, TITLE: Learning to Express in Knowledge-Grounded Conversation
<https://aclanthology.org/2022.naacl-main.164>
AUTHORS: Xueliang Zhao, Tingchen Fu, Chongyang Tao, Wei Wu, Dongyan Zhao, Rui Yan
HIGHLIGHT: In this work, we mainly consider two aspects of knowledge expression, namely the structure of the response and style of the content in each part.
- 165, TITLE: End-to-End Chinese Speaker Identification
<https://aclanthology.org/2022.naacl-main.165>
AUTHORS: Dian Yu, Ben Zhou, Dong Yu
HIGHLIGHT: To make large end-to-end models possible, we design a new annotation guideline that regards SI as span extraction from the local context, and we annotate by far the largest SI dataset for Chinese named CSI based on eighteen novels.
- 166, TITLE: MINION: a Large-Scale and Diverse Dataset for Multilingual Event Detection
<https://aclanthology.org/2022.naacl-main.166>
AUTHORS: Amir Pouran Ben Veysch, Minh Van Nguyen, Franck Dernoncourt, Thien Nguyen
HIGHLIGHT: In addition, the current datasets are often small and not accessible to the public. To overcome those shortcomings, we introduce a new large-scale multilingual dataset for ED (called MINION) that consistently annotates events for 8 different languages; 5 of them have not been supported by existing multilingual datasets.
- 167, TITLE: Do Prompt-Based Models Really Understand the Meaning of Their Prompts?
<https://aclanthology.org/2022.naacl-main.167>
AUTHORS: Albert Webson, Ellie Pavlick
HIGHLIGHT: In this study, we experiment with over 30 prompts manually written for natural language inference (NLI).
- 168, TITLE: GPL: Generative Pseudo Labeling for Unsupervised Domain Adaptation of Dense Retrieval
<https://aclanthology.org/2022.naacl-main.168>
AUTHORS: Kexin Wang, Nandan Thakur, Nils Reimers, Iryna Gurevych
HIGHLIGHT: In this paper, we propose the novel unsupervised domain adaptation method Generative Pseudo Labeling (GPL), which combines a query generator with pseudo labeling from a cross-encoder.
- 169, TITLE: Sparse Distillation: Speeding Up Text Classification by Using Bigger Student Models
<https://aclanthology.org/2022.naacl-main.169>
AUTHORS: Qinyuan Ye, Madian Khabsa, Mike Lewis, Sinong Wang, Xiang Ren, Aaron Jaech
HIGHLIGHT: Therefore, the improved inference speed may still be unsatisfactory for real-time or high-volume use cases. In this paper, we aim to further push the limit of inference speed by distilling teacher models into bigger, sparser student models - bigger in that they scale up to billions of parameters; sparser in that most of the model parameters are n-gram embeddings.
- 170, TITLE: Towards Understanding Large-Scale Discourse Structures in Pre-Trained and Fine-Tuned Language Models
<https://aclanthology.org/2022.naacl-main.170>
AUTHORS: Patrick Huber, Giuseppe Carenini
HIGHLIGHT: In this paper, we extend the line of BERTology work by focusing on the important, yet less explored, alignment of pre-trained and fine-tuned PLMs with large-scale discourse structures.
- 171, TITLE: SAIS: Supervising and Augmenting Intermediate Steps for Document-Level Relation Extraction
<https://aclanthology.org/2022.naacl-main.171>
AUTHORS: Yuxin Xiao, Zecheng Zhang, Yuning Mao, Carl Yang, Jiawei Han
HIGHLIGHT: In contrast, we propose to explicitly teach the model to capture relevant contexts and entity types by supervising and augmenting intermediate steps (SAIS) for RE.
- 172, TITLE: LITE: Intent-based Task Representation Learning Using Weak Supervision
<https://aclanthology.org/2022.naacl-main.172>
AUTHORS: Naoki Otani, Michael Gamon, Sujay Kumar Jauhar, Mei Yang, Sri Raghu Malireddi, Oriana Riva
HIGHLIGHT: Yet, understanding and representing their meaning is the first step towards providing intelligent assistance for to-do management. We address this problem by proposing a neural multi-task learning framework, LITE, which extracts representations of English to-do tasks with a multi-head attention mechanism on top of a pre-trained text encoder.
- 173, TITLE: Does Summary Evaluation Survive Translation to Other Languages?
<https://aclanthology.org/2022.naacl-main.173>
AUTHORS: Spencer Braun, Oleg Vasilyev, Neslihan Iskender, John Bohannon
HIGHLIGHT: To investigate how much we can trust machine translation of summarization datasets, we translate the English SummEval dataset to seven languages and compare performances across automatic evaluation measures.
- 174, TITLE: A Shoulder to Cry on: Towards A Motivational Virtual Assistant for Assuaging Mental Agony

<https://aclanthology.org/2022.naacl-main.174>

AUTHORS: Tulika Saha, Saichethan Reddy, Anindya Das, Sriparna Saha, Pushpak Bhattacharyya
HIGHLIGHT: In this paper, we propose a VA that can act as the first point of contact and comfort for mental health patients.

175, TITLE: SueNes: A Weakly Supervised Approach to Evaluating Single-Document Summarization via Negative Sampling

<https://aclanthology.org/2022.naacl-main.175>

AUTHORS: Forrest Bao, Ge Luo, Hebi Li, Minghui Qiu, Yinfei Yang, Youbiao He, Cen Chen
HIGHLIGHT: In this paper, we present a proof-of-concept study to a weakly supervised summary evaluation approach without the presence of reference summaries.

176, TITLE: Combating the Curse of Multilinguality in Cross-Lingual WSD by Aligning Sparse Contextualized Word Representations

<https://aclanthology.org/2022.naacl-main.176>

AUTHORS: G?bor Berend
HIGHLIGHT: In this paper, we advocate for using large pre-trained monolingual language models in cross lingual zero-shot word sense disambiguation (WSD) coupled with a contextualized mapping mechanism.

177, TITLE: Cheat Codes to Quantify Missing Source Information in Neural Machine Translation

<https://aclanthology.org/2022.naacl-main.177>

AUTHORS: Proyag Pal, Kenneth Heafield
HIGHLIGHT: This paper describes a method to quantify the amount of information $H(t|s)$ added by the target sentence t that is not present in the source s in a neural machine translation system.

178, TITLE: WiC = TSV = WSD: On the Equivalence of Three Semantic Tasks

<https://aclanthology.org/2022.naacl-main.178>

AUTHORS: Bradley Hauer, Grzegorz Kondrak
HIGHLIGHT: In this paper, we establish the exact relationship between WiC and WSD, as well as the related task of target sense verification (TSV).

179, TITLE: What do tokens know about their characters and how do they know it?

<https://aclanthology.org/2022.naacl-main.179>

AUTHORS: Ayush Kaushal, Kyle Mahowald
HIGHLIGHT: Here, studying a range of models (e.g., GPT- J, BERT, RoBERTa, GloVe), we probe what word pieces encode about character-level information by training classifiers to predict the presence or absence of a particular alphabetical character in a token, based on its embedding (e.g., probing whether the model embedding for "cat" encodes that it contains the character "a").

180, TITLE: AnswerSumm: A Manually-Curated Dataset and Pipeline for Answer Summarization

<https://aclanthology.org/2022.naacl-main.180>

AUTHORS: Alexander Fabbri, Xiaojian Wu, Sridi Iyer, Haoran Li, Mona Diab
HIGHLIGHT: This work introduces a novel dataset of 4,631 CQA threads for answer summarization curated by professional linguists.

181, TITLE: Paragraph-based Transformer Pre-training for Multi-Sentence Inference

<https://aclanthology.org/2022.naacl-main.181>

AUTHORS: Luca Di Lello, Siddhant Garg, Luca Soldaini, Alessandro Moschitti
HIGHLIGHT: In this paper, we first show that popular pre-trained transformers perform poorly when used for fine-tuning on multi-candidate inference tasks. We then propose a new pre-training objective that models the paragraph-level semantics across multiple input sentences.

182, TITLE: Text Style Transfer via Optimal Transport

<https://aclanthology.org/2022.naacl-main.182>

AUTHORS: Nasim Nouri
HIGHLIGHT: In this work, we propose a novel method based on Optimal Transport for TST to simultaneously incorporate syntactic and semantic information into similarity computation between the source and the converted text.

183, TITLE: Exploring the Role of Task Transferability in Large-Scale Multi-Task Learning

<https://aclanthology.org/2022.naacl-main.183>

AUTHORS: Vishakh Padmakumar, Leonard Lausen, Miguel Ballesteros, Sheng Zha, He He, George Karypis
HIGHLIGHT: In this work, we aim to disentangle the effect of scale and relatedness of tasks in multi-task representation learning.

184, TITLE: Interactive Query-Assisted Summarization via Deep Reinforcement Learning

<https://aclanthology.org/2022.naacl-main.184>

AUTHORS: Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Ido Dagan, Yael Amsterdamer
HIGHLIGHT: To that end, we propose two novel deep reinforcement learning models for the task that address, respectively, the subtask of summarizing salient information that adheres to user queries, and the subtask of listing suggested queries to assist users throughout their exploration.

185, TITLE: Data Augmentation with Dual Training for Offensive Span Detection

<https://aclanthology.org/2022.naacl-main.185>

AUTHORS: Nasim Nouri

HIGHLIGHT: One of the challenges to train a model for this novel setting is the lack of enough training data. To address this limitation, in this work we propose a novel method in which the large-scale pre-trained language model GPT-2 is employed to generate synthetic training data for OSD.

186, TITLE: Training Mixed-Domain Translation Models via Federated Learning

<https://aclanthology.org/2022.naacl-main.186>

AUTHORS: Peyman Passban, Tanya Roosta, Rahul Gupta, Ankit Chadha, Clement Chung

HIGHLIGHT: In this work, we leverage federated learning (FL) in order to tackle the problem.

187, TITLE: QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization

<https://aclanthology.org/2022.naacl-main.187>

AUTHORS: Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu, Caiming Xiong

HIGHLIGHT: In this work, we conduct an extensive comparison of entailment and QA-based metrics, demonstrating that carefully choosing the components of a QA-based metric, especially question generation and answerability classification, is critical to performance.

188, TITLE: How Gender Debiasing Affects Internal Model Representations, and Why It Matters

<https://aclanthology.org/2022.naacl-main.188>

AUTHORS: Hadas Orgad, Seraphina Goldfarb-Tarrant, Yonatan Belinkov

HIGHLIGHT: In this work, we illuminate this relationship by measuring both quantities together: we debias a model during downstream fine-tuning, which reduces extrinsic bias, and measure the effect on intrinsic bias, which is operationalized as bias extractability with information-theoretic probing.

189, TITLE: A Structured Span Selector

<https://aclanthology.org/2022.naacl-main.189>

AUTHORS: Tianyu Liu, Yuchen Jiang, Ryan Cotterell, Mrinmaya Sachan

HIGHLIGHT: In this paper, we propose a novel grammar-based structured span selection model which learns to make use of the partial span-level annotation provided for such problems.

190, TITLE: Unified Semantic Typing with Meaningful Label Inference

<https://aclanthology.org/2022.naacl-main.190>

AUTHORS: James Y. Huang, Bangzheng Li, Jiashu Xu, Muhao Chen

HIGHLIGHT: In this paper, we present UniST, a unified framework for semantic typing that captures label semantics by projecting both inputs and labels into a joint semantic embedding space.

191, TITLE: Learning To Retrieve Prompts for In-Context Learning

<https://aclanthology.org/2022.naacl-main.191>

AUTHORS: Ohad Rubin, Jonathan Herzig, Jonathan Berant

HIGHLIGHT: In this work, we propose an efficient method for retrieving prompts for in-context learning using annotated data and an LM.

192, TITLE: Necessity and Sufficiency for Explaining Text Classifiers: A Case Study in Hate Speech Detection

<https://aclanthology.org/2022.naacl-main.192>

AUTHORS: Esma Balkir, Isar Nejadgholi, Kathleen Fraser, Svetlana Kiritchenko

HIGHLIGHT: We present a novel feature attribution method for explaining text classifiers, and analyze it in the context of hate speech detection.

193, TITLE: Learning to Retrieve Passages without Supervision

<https://aclanthology.org/2022.naacl-main.193>

AUTHORS: Ori Ram, Gal Shachaf, Omer Levy, Jonathan Berant, Amir Globerson

HIGHLIGHT: Dense retrievers for open-domain question answering (ODQA) have been shown to achieve impressive performance by training on large datasets of question-passage pairs. In this work we ask whether this dependence on labeled data can be reduced via unsupervised pretraining that is geared towards ODQA.

194, TITLE: Re2G: Retrieve, Rerank, Generate

<https://aclanthology.org/2022.naacl-main.194>

AUTHORS: Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Naik, Pengshan Cai, Alfio Gliozzo

HIGHLIGHT: We build on this line of research, proposing Re2G, which combines both neural initial retrieval and reranking into a BART-based sequence-to-sequence generation.

195, TITLE: Don't sweat the small stuff, classify the rest: Sample Shielding to protect text classifiers against adversarial attacks

<https://aclanthology.org/2022.naacl-main.195>

AUTHORS: Jonathan Rusert, Padmini Srinivasan

HIGHLIGHT: We shield three popular DL text classifiers with Sample Shielding, test their resilience against four SOTA attackers across three datasets in a realistic threat setting.

196, TITLE: Federated Learning with Noisy User Feedback

<https://aclanthology.org/2022.naacl-main.196>

AUTHORS: Rahul Sharma, Anil Ramakrishna, Ansel MacLaughlin, Anna Rumshisky, Jimit Majmudar, Clement Chung, Salman Avestimehr, Rahul Gupta
HIGHLIGHT: In this work, we propose a strategy for training FL models using positive and negative user feedback.

197, TITLE: Gender Bias in Masked Language Models for Multiple Languages

<https://aclanthology.org/2022.naacl-main.197>

AUTHORS: Masahiro Kaneko, Aizhan Imankulova, Danushka Bollegala, Naoaki Okazaki
HIGHLIGHT: We propose Multilingual Bias Evaluation (MBE) score, to evaluate bias in various languages using only English attribute word lists and parallel corpora between the target language and English without requiring manually annotated data.

198, TITLE: Multi-Domain Targeted Sentiment Analysis

<https://aclanthology.org/2022.naacl-main.198>

AUTHORS: Orith Toledo-Ronen, Matan Orbach, Yoav Katz, Noam Slonim
HIGHLIGHT: To address this scenario, we present a multi-domain TSA system based on augmenting a given training set with diverse weak labels from assorted domains.

199, TITLE: Falsesum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization

<https://aclanthology.org/2022.naacl-main.199>

AUTHORS: Prasetya Utama, Joshua Bambrick, Nafise Moosavi, Iryna Gurevych
HIGHLIGHT: However, state-of-the-art NLI models perform poorly in this context due to their inability to generalize to the target task. In this work, we show that NLI models can be effective for this task when the training data is augmented with high-quality task-oriented examples.

200, TITLE: Dynamic Gazetteer Integration in Multilingual Models for Cross-Lingual and Cross-Domain Named Entity Recognition

<https://aclanthology.org/2022.naacl-main.200>

AUTHORS: Besnik Fetahu, Anjie Fang, Oleg Rokhlenko, Shervin Malmasi
HIGHLIGHT: We propose an approach that with limited effort and data, addresses the NER knowledge gap across languages and domains.

201, TITLE: MetaICL: Learning to Learn In Context

<https://aclanthology.org/2022.naacl-main.201>

AUTHORS: Sewon Min, Mike Lewis, Luke Zettlemoyer, Hannaneh Hajishirzi
HIGHLIGHT: We introduce MetaICL (Meta-training for In-Context Learning), a new meta-training framework for few-shot learning where a pretrained language model is tuned to do in-context learning on a large set of training tasks.

202, TITLE: Enhancing Knowledge Selection for Grounded Dialogues via Document Semantic Graphs

<https://aclanthology.org/2022.naacl-main.202>

AUTHORS: Sha Li, Mahdi Namazifar, Di Jin, Mohit Bansal, Heng Ji, Yang Liu, Dilek Hakkani-Tur
HIGHLIGHT: Existing models treat knowledge selection as a sentence ranking or classification problem where each sentence is handled individually, ignoring the internal semantic connection between sentences. In this work, we propose to automatically convert the background knowledge documents into document semantic graphs and then perform knowledge selection over such graphs.

203, TITLE: Using Natural Sentence Prompts for Understanding Biases in Language Models

<https://aclanthology.org/2022.naacl-main.203>

AUTHORS: Sarah Alnegheimish, Alicia Guo, Yi Sun
HIGHLIGHT: This dependence traces back to the need of prompt-style dataset to trigger specific behaviors of language models. In this paper, we address this gap by creating a prompt dataset with respect to occupations collected from real-world natural sentences present in Wikipedia.

204, TITLE: Robust Conversational Agents against Imperceptible Toxicity Triggers

<https://aclanthology.org/2022.naacl-main.204>

AUTHORS: Ninareh Mehrabi, Ahmad Beirami, Fred Morstatter, Aram Galstyan
HIGHLIGHT: In this work, we propose attacks against conversational agents that are imperceptible, i.e., they fit the conversation in terms of coherency, relevancy, and fluency, while they are effective and scalable, i.e., they can automatically trigger the system into generating toxic language.

205, TITLE: Selective Differential Privacy for Language Modeling

<https://aclanthology.org/2022.naacl-main.205>

AUTHORS: Weiyang Shi, Aiqi Cui, Evan Li, Ruoxi Jia, Zhou Yu
HIGHLIGHT: Given that the private information in natural language is sparse (for example, the bulk of an email might not carry personally identifiable information), we propose a new privacy notion, selective differential privacy, to provide rigorous privacy guarantees on the sensitive portion of the data to improve model utility.

206, TITLE: Do Trajectories Encode Verb Meaning?

<https://aclanthology.org/2022.naacl-main.206>

AUTHORS: Dylan Ebert, Chen Sun, Ellie Pavlick

HIGHLIGHT: In this paper, we investigate the extent to which trajectories (i.e. the position and rotation of objects over time) naturally encode verb semantics.

207, **TITLE:** Long Context Question Answering via Supervised Contrastive Learning

<https://aclanthology.org/2022.naacl-main.207>

AUTHORS: Avi Caciularu, Ido Dagan, Jacob Goldberger, Arman Cohan

HIGHLIGHT: In this work, we propose a novel method for equipping long-context QA models with an additional sequence-level objective for better identification of the supporting evidence.

208, **TITLE:** The USMLE Step 2 Clinical Skills Patient Note Corpus

<https://aclanthology.org/2022.naacl-main.208>

AUTHORS: Victoria Yaneva, Janet Mee, Le Ha, Polina Harik, Michael Jodoin, Alex Mehaber

HIGHLIGHT: This paper presents a corpus of 43,985 clinical patient notes (PNs) written by 35,156 examinees during the high-stakes USMLE Step 2 Clinical Skills examination.

209, **TITLE:** Learning to Borrow- Relation Representation for Without-Mention Entity-Pairs for Knowledge Graph Completion

<https://aclanthology.org/2022.naacl-main.209>

AUTHORS: Huda Hakami, Mona Hakami, Angrosh Mandya, Danushka Bollegala

HIGHLIGHT: We propose a supervised borrowing method, SuperBorrow, that learns to score the suitability of an LDP to represent a without-mentions entity pair using pre-trained entity embeddings and contextualised LDP representations.

210, **TITLE:** Improving Entity Disambiguation by Reasoning over a Knowledge Base

<https://aclanthology.org/2022.naacl-main.210>

AUTHORS: Tom Ayoola, Joseph Fisher, Andrea Pierleoni

HIGHLIGHT: To allow the use of all KB facts, as well as descriptions and types, we introduce an ED model which links entities by reasoning over a symbolic knowledge base in a fully differentiable fashion.

211, **TITLE:** Modal Dependency Parsing via Language Model Priming

<https://aclanthology.org/2022.naacl-main.211>

AUTHORS: Jiarui Yao, Nianwen Xue, Bonan Min

HIGHLIGHT: We design a modal dependency parser that is based on priming pre-trained language models, and evaluate the parser on two data sets.

212, **TITLE:** Document-Level Relation Extraction with Sentences Importance Estimation and Focusing

<https://aclanthology.org/2022.naacl-main.212>

AUTHORS: Wang Xu, Kehai Chen, Lili Mou, Tiejun Zhao

HIGHLIGHT: To this end, we propose a Sentence Importance Estimation and Focusing (SIEF) framework for DocRE, where we design a sentence importance score and a sentence focusing loss, encouraging DocRE models to focus on evidence sentences.

213, **TITLE:** Are All the Datasets in Benchmark Necessary? A Pilot Study of Dataset Evaluation for Text Classification

<https://aclanthology.org/2022.naacl-main.213>

AUTHORS: Yang Xiao, Jinlan Fu, See-Kiong Ng, Pengfei Liu

HIGHLIGHT: In this paper, we ask the research question of whether all the datasets in the benchmark are necessary.

214, **TITLE:** Triggerless Backdoor Attack for NLP Tasks with Clean Labels

<https://aclanthology.org/2022.naacl-main.214>

AUTHORS: Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, Yi Yang, Shangwei Guo, Chun Fan

HIGHLIGHT: To generate poisoned clean-labeled examples, we propose a sentence generation model based on the genetic algorithm to cater to the non-differentiable characteristic of text data.

215, **TITLE:** PPL-MCTS: Constrained Textual Generation Through Discriminator-Guided MCTS Decoding

<https://aclanthology.org/2022.naacl-main.215>

AUTHORS: Antoine Chaffin, Vincent Claveau, Ewa Kijak

HIGHLIGHT: Large language models (LM) based on Transformers allow to generate plausible long texts. In this paper, we explore how this generation can be further controlled at decoding time to satisfy certain constraints (e.g. being non-toxic, conveying certain emotions, using a specific writing style, etc.) without fine-tuning the LM.

216, **TITLE:** Interpretable Proof Generation via Iterative Backward Reasoning

<https://aclanthology.org/2022.naacl-main.216>

AUTHORS: Hanhao Qu, Yu Cao, Jun Gao, Liang Ding, Ruifeng Xu

HIGHLIGHT: We present IBR, an Iterative Backward Reasoning model to solve the proof generation tasks on rule-based Question Answering (QA), where models are required to reason over a series of textual rules and facts to find out the related proof path and derive the final answer.

217, **TITLE:** Domain Confused Contrastive Learning for Unsupervised Domain Adaptation

<https://aclanthology.org/2022.naacl-main.217>

AUTHORS: Quanyu Long, Tianze Luo, Wenya Wang, Sinno Pan

HIGHLIGHT: In this work, we study Unsupervised Domain Adaptation (UDA) in a challenging self-supervised approach.

- 218, TITLE: Incorporating Centering Theory into Neural Coreference Resolution
<https://aclanthology.org/2022.naacl-main.218>
AUTHORS: Haixia Chai, Michael Strube
HIGHLIGHT: In this paper, we propose to incorporate centering transitions derived from centering theory in the form of a graph into a neural coreference model.
- 219, TITLE: Progressive Class Semantic Matching for Semi-supervised Text Classification
<https://aclanthology.org/2022.naacl-main.219>
AUTHORS: Haiming Xu, Lingqiao Liu, Ehsan Abbasnejad
HIGHLIGHT: Specifically, we propose a joint semi-supervised learning process that can progressively build a standard K-way classifier and a matching network for the input text and the Class Semantic Representation (CSR).
- 220, TITLE: Low Resource Style Transfer via Domain Adaptive Meta Learning
<https://aclanthology.org/2022.naacl-main.220>
AUTHORS: Xiangyang Li, Xiang Long, Yu Xia, Sujian Li
HIGHLIGHT: In this work, we propose DAML-ATM (Domain Adaptive Meta-Learning with Adversarial Transfer Model), which consists of two parts: DAML and ATM.
- 221, TITLE: Features or Spurious Artifacts? Data-centric Baselines for Fair and Robust Hate Speech Detection
<https://aclanthology.org/2022.naacl-main.221>
AUTHORS: Alan Ramponi, Sara Tonelli
HIGHLIGHT: In this paper we critically analyze lexical biases in hate speech detection via a cross-platform study, disentangling various types of spurious and authentic artifacts and analyzing their impact on out-of-distribution fairness and robustness.
- 222, TITLE: Document-Level Event Argument Extraction by Leveraging Redundant Information and Closed Boundary Loss
<https://aclanthology.org/2022.naacl-main.222>
AUTHORS: Hanzhang Zhou, Kezhi Mao
HIGHLIGHT: In this paper, to make use of redundant event information underlying a document, we build an entity coreference graph with the graph2token module to produce a comprehensive and coreference-aware representation for every entity and then build an entity summary graph to merge the multiple extraction results.
- 223, TITLE: A Few Thousand Translations Go a Long Way! Leveraging Pre-trained Models for African News Translation
<https://aclanthology.org/2022.naacl-main.223>
AUTHORS: David Adelani, Jesujoba Alabi, Angela Fan, Julia Kreutzer, Xiaoyu Shen, Machel Reid, Dana Ruiter, Dietrich Klakow, Peter Nabende, Ernie Chang, Tajuddeen Gwadabe, Freshia Sackey, Bonaventure F. P. Dossou, Chris Emezue, Colin Leong, Michael Beukman, Shamsuddeen Muhammad, Guyo Jarso, Oreen Yousuf, Andre Niyongabo Rubungo, Gilles Hacheme, Eric Peter Wairagala, Muhammad Umair Nasir, Benjamin Ajibade, Tunde Ajayi, Yvonne Gitau, Jade Abbott, Mohamed Ahmed, Millicent Ochieng, Anuoluwapo Aremu, Perez Ogayo, Jonathan Mukiibi, Fatoumata Ouoba Kabore, Godson Kalipe, Derguene Mbaye, Allahsera Auguste Tapo, Victoire Memdjokam Kogagne, Edwin Munkoh-Buabeng, Valencia Wagner, Idris Abdulmumin, Ayodele Awokoya, Happy Buzaaba, Blessing Sibanda, Andiswa Bukula, Sam Manthalu
HIGHLIGHT: This work investigates how to optimally leverage existing pre-trained models to create low-resource translation systems for 16 African languages.
- 224, TITLE: Should We Rely on Entity Mentions for Relation Extraction? Debiasing Relation Extraction with Counterfactual Analysis
<https://aclanthology.org/2022.naacl-main.224>
AUTHORS: Yiwei Wang, Muhao Chen, Wenxuan Zhou, Yujun Cai, Yuxuan Liang, Dayiheng Liu, Baosong Yang, Juncheng Liu, Bryan Hooi
HIGHLIGHT: In this paper, we propose the CoRE (Counterfactual Analysis based Relation Extraction) debiasing method that guides the RE models to focus on the main effects of textual context without losing the entity information.
- 225, TITLE: Analyzing Encoded Concepts in Transformer Language Models
<https://aclanthology.org/2022.naacl-main.225>
AUTHORS: Hassan Sajjad, Nadir Durrani, Fahim Dalvi, Firoj Alam, Abdul Khan, Jia Xu
HIGHLIGHT: We propose a novel framework ConceptX, to analyze how latent concepts are encoded in representations learned within pre-trained lan-guage models.
- 226, TITLE: Boosted Dense Retriever
<https://aclanthology.org/2022.naacl-main.226>
AUTHORS: Patrick Lewis, Barlas Oguz, Wenhan Xiong, Fabio Petroni, Scott Yih, Sebastian Riedel
HIGHLIGHT: We propose DrBoost, a dense retrieval ensemble inspired by boosting.
- 227, TITLE: MuCGEC: a Multi-Reference Multi-Source Evaluation Dataset for Chinese Grammatical Error Correction
<https://aclanthology.org/2022.naacl-main.227>
AUTHORS: Yue Zhang, Zhenghua Li, Zuyi Bao, Jiacheng Li, Bo Zhang, Chen Li, Fei Huang, Min Zhang
HIGHLIGHT: This paper presents MuCGEC, a multi-reference multi-source evaluation dataset for Chinese Grammatical Error Correction (CGEC), consisting of 7,063 sentences collected from three Chinese-as-a-Second-Language (CSL) learner sources.
- 228, TITLE: NeuS: Neutral Multi-News Summarization for Mitigating Framing Bias

- <https://aclanthology.org/2022.naacl-main.228>
AUTHORS: Nayeon Lee, Yejin Bang, Tiezheng Yu, Andrea Madotto, Pascale Fung
HIGHLIGHT: We propose a new task, a neutral summary generation from multiple news articles of the varying political leaningsto facilitate balanced and unbiased news reading.
- 229, TITLE: Enhance Incomplete Utterance Restoration by Joint Learning Token Extraction and Text Generation
<https://aclanthology.org/2022.naacl-main.229>
AUTHORS: Shumpei Inoue, Tsungwei Liu, Son Nguyen, Minh-Tien Nguyen
HIGHLIGHT: This paper introduces a model for incomplete utterance restoration (IUR) called JET (Joint learning token Extraction and Text generation).
- 230, TITLE: Efficient Constituency Tree based Encoding for Natural Language to Bash Translation
<https://aclanthology.org/2022.naacl-main.230>
AUTHORS: Shikhar Bharadwaj, Shirish Shevade
HIGHLIGHT: Recent works on natural language to Bash translation have made significant advances, but none of the previous methods utilize the problem's inherent structure. We identify this structure andpropose a Segmented Invocation Transformer (SIT) that utilizes the information from the constituency parse tree of the natural language text.
- 231, TITLE: Privacy-Preserving Text Classification on BERT Embeddings with Homomorphic Encryption
<https://aclanthology.org/2022.naacl-main.231>
AUTHORS: Garam Lee, Minsoo Kim, Jai Hyun Park, Seung-won Hwang, Jung Hee Cheon
HIGHLIGHT: However, recent research has shown that embeddings can potentially leak private information about sensitive attributes of the text, and in some cases, can be inverted to recover the original input text. To address these growing privacy challenges, we propose a privatization mechanism for embeddings based on homomorphic encryption, to prevent potential leakage of any piece of information in the process of text classification.
- 232, TITLE: ITA: Image-Text Alignments for Multi-Modal Named Entity Recognition
<https://aclanthology.org/2022.naacl-main.232>
AUTHORS: Xinyu Wang, Min Gui, Yong Jiang, Zixia Jia, Nguyen Bach, Tao Wang, Zhongqiang Huang, Kewei Tu
HIGHLIGHT: As text representations take the most important role in MNER, in this paper, we propose Image-text Alignments (ITA) to align image features into the textual space, so that the attention mechanism in transformer-based pretrained textual embeddings can be better utilized.
- 233, TITLE: A Dataset for N-ary Relation Extraction of Drug Combinations
<https://aclanthology.org/2022.naacl-main.233>
AUTHORS: Aryeh Tiktinsky, Vijay Viswanathan, Danna Niezni, Dana Meron Azagury, Yosi Shamay, Hillel Taub-Tabib, Tom Hope, Yoav Goldberg
HIGHLIGHT: To assist medical professionals in identifying beneficial drug-combinations, we construct an expert-annotated dataset for extracting information about the efficacy of drug combinations from the scientific literature.
- 234, TITLE: Curriculum: A Broad-Coverage Benchmark for Linguistic Phenomena in Natural Language Understanding
<https://aclanthology.org/2022.naacl-main.234>
AUTHORS: Zeming Chen, Qiyue Gao
HIGHLIGHT: In this paper, we introduce Curriculum as a new format of NLI benchmark for evaluation of broad-coverage linguistic phenomena.
- 235, TITLE: Neural Language Taskonomy: Which NLP Tasks are the most Predictive of fMRI Brain Activity?
<https://aclanthology.org/2022.naacl-main.235>
AUTHORS: Subba Reddy Oota, Jashn Arora, Veeral Agarwal, Mounika Marreddy, Manish Gupta, Bapi Surampudi
HIGHLIGHT: In this work, we explore transfer learning from representations learned for ten popular natural language processing tasks (two syntactic and eight semantic) for predicting brain responses from two diverse datasets: Pereira (subjects reading sentences from paragraphs) and Narratives (subjects listening to the spoken stories).
- 236, TITLE: FactGraph: Evaluating Factuality in Summarization with Semantic Graph Representations
<https://aclanthology.org/2022.naacl-main.236>
AUTHORS: Leonardo Ribeiro, Mengwen Liu, Iryna Gurevych, Markus Dreyer, Mohit Bansal
HIGHLIGHT: Recent works have shown promising improvements in factuality error identification using text or dependency arc entailments; however, they do not consider the entire semantic graph simultaneously. To this end, we propose FactGraph, a method that decomposes the document and the summary into structured meaning representations (MR), which are more suitable for factuality evaluation.
- 237, TITLE: Unsupervised Paraphrasability Prediction for Compound Nominalizations
<https://aclanthology.org/2022.naacl-main.237>
AUTHORS: John Sie Yuen Lee, Ho Hung Lim, Carol Carol Webster
HIGHLIGHT: This paper investigates unsupervised prediction of paraphrasability, which determines whether the prenominal modifier of a nominalization can be re-written as a noun or adverb in a clausal paraphrase.
- 238, TITLE: Global Entity Disambiguation with BERT
<https://aclanthology.org/2022.naacl-main.238>
AUTHORS: Ikuya Yamada, Koki Washio, Hiroyuki Shindo, Yuji Matsumoto

- HIGHLIGHT:** We propose a global entity disambiguation (ED) model based on BERT.
- 239, **TITLE:** Clues Before Answers: Generation-Enhanced Multiple-Choice QA
<https://aclanthology.org/2022.naacl-main.239>
AUTHORS: Zixian Huang, Ao Wu, Jiaying Zhou, Yu Gu, Yue Zhao, Gong Cheng
HIGHLIGHT: To exploit the generation capability and underlying knowledge of a pre-trained encoder-decoder model, in this paper, we propose a generation-enhanced MCQA model named GenMC.
- 240, **TITLE:** Towards Efficient NLP: A Standard Evaluation and A Strong Baseline
<https://aclanthology.org/2022.naacl-main.240>
AUTHORS: Xiangyang Liu, Tianxiang Sun, Junliang He, Jiawen Wu, Lingling Wu, Xinyu Zhang, Hao Jiang, Zhao Cao, Xuanjing Huang, Xipeng Qiu
HIGHLIGHT: To that end, this work presents ELUE (Efficient Language Understanding Evaluation), a standard evaluation, and a public leaderboard for efficient NLP models.
- 241, **TITLE:** Stylized Knowledge-Grounded Dialogue Generation via Disentangled Template Rewriting
<https://aclanthology.org/2022.naacl-main.241>
AUTHORS: Qingfeng Sun, Can Xu, Huang Hu, Yujing Wang, Jian Miao, Xiubo Geng, Yining Chen, Fei Xu, Daxin Jiang
HIGHLIGHT: In this paper, we propose a novel disentangled template rewriting (DTR) method which generates responses via combing disentangled style templates (from monolingual stylized corpus) and content templates (from KDG corpus).
- 242, **TITLE:** LUNA: Learning Slot-Turn Alignment for Dialogue State Tracking
<https://aclanthology.org/2022.naacl-main.242>
AUTHORS: Yifan Wang, Jing Zhao, Junwei Bao, Chaoqun Duan, Youzheng Wu, Xiaodong He
HIGHLIGHT: This could lead to suboptimal results due to the information introduced from irrelevant utterances in the dialogue history, which may be useless and can even cause confusion. To address this problem, we propose LUNA, a SLOT-TURN Alignment enhanced approach.
- 243, **TITLE:** Crossroads, Buildings and Neighborhoods: A Dataset for Fine-grained Location Recognition
<https://aclanthology.org/2022.naacl-main.243>
AUTHORS: Pei Chen, Haotian Xu, Cheng Zhang, Ruihong Huang
HIGHLIGHT: In this paper, we introduce a new dataset HarveyNER with fine-grained locations annotated in tweets.
- 244, **TITLE:** Tricks for Training Sparse Translation Models
<https://aclanthology.org/2022.naacl-main.244>
AUTHORS: Dheeru Dua, Shruti Bhosale, Vedanuj Goswami, James Cross, Mike Lewis, Angela Fan
HIGHLIGHT: We find that that sparse architectures for multilingual machine translation can perform poorly out of the box and propose two straightforward techniques to mitigate this - a temperature heating mechanism and dense pre-training.
- 245, **TITLE:** Persona-Guided Planning for Controlling the Protagonist's Persona in Story Generation
<https://aclanthology.org/2022.naacl-main.245>
AUTHORS: Zhixin Zhang, Jiaxin Wen, Jian Guan, Minlie Huang
HIGHLIGHT: In this paper, we aim to control the protagonist's persona in story generation, i.e., generating a story from a leading context and a persona description, where the protagonist should exhibit the specified personality through a coherent event sequence.
- 246, **TITLE:** CHEF: A Pilot Chinese Dataset for Evidence-Based Fact-Checking
<https://aclanthology.org/2022.naacl-main.246>
AUTHORS: Xuming Hu, Zhijiang Guo, GuanYu Wu, Aiwei Liu, Lijie Wen, Philip Yu
HIGHLIGHT: Datasets and tools available in other languages, such as Chinese, are limited. In order to bridge this gap, we construct CHEF, the first CHINESE Evidence-based Fact-checking dataset of 10K real-world claims.
- 247, **TITLE:** VGNMN: Video-grounded Neural Module Networks for Video-Grounded Dialogue Systems
<https://aclanthology.org/2022.naacl-main.247>
AUTHORS: Hung Le, Nancy Chen, Steven Hoi
HIGHLIGHT: Motivated by recent NMN approaches on image-grounded tasks, we introduce Video-grounded Neural Module Network (VGNMN) to model the information retrieval process in video-grounded language tasks as a pipeline of neural modules.
- 248, **TITLE:** Multimodal Dialogue State Tracking
<https://aclanthology.org/2022.naacl-main.248>
AUTHORS: Hung Le, Nancy Chen, Steven Hoi
HIGHLIGHT: In this paper, we propose to extend the definition of dialogue state tracking to multimodality.
- 249, **TITLE:** On the Use of Bert for Automated Essay Scoring: Joint Learning of Multi-Scale Essay Representation
<https://aclanthology.org/2022.naacl-main.249>
AUTHORS: Yongjie Wang, Chuang Wang, Ruobing Li, Hui Lin
HIGHLIGHT: In this paper, we introduce a novel multi-scale essay representation for BERT that can be jointly learned.
- 250, **TITLE:** Recognition of They/Them as Singular Personal Pronouns in Coreference Resolution
<https://aclanthology.org/2022.naacl-main.250>

- AUTHORS: Connor Baumler, Rachel Rudinger
HIGHLIGHT: We introduce a new benchmark for coreference resolution systems which evaluates singular personal "they" recognition.
- 251, TITLE: TWEETSPIN: Fine-grained Propaganda Detection in Social Media Using Multi-View Representations
<https://aclanthology.org/2022.naacl-main.251>
AUTHORS: Prashanth Vijayaraghavan, Soroush Vosoughi
HIGHLIGHT: In this work, we present TWEETSPIN, a dataset containing tweets that are weakly annotated with different fine-grained propaganda techniques, and propose a neural approach to detect and categorize propaganda tweets across those fine-grained categories.
- 252, TITLE: UserIdentifier: Implicit User Representations for Simple and Effective Personalized Sentiment Analysis
<https://aclanthology.org/2022.naacl-main.252>
AUTHORS: Fatemehsadat Mireshghallah, Vaishnavi Shrivastava, Milad Shokouhi, Taylor Berg-Kirkpatrick, Robert Sim, Dimitrios Dimitriadis
HIGHLIGHT: Contrary to widely-used personalization techniques based on few-shot and meta-learning, we propose UserIdentifier, a novel scheme for training a single shared model for all users.
- 253, TITLE: Improving Neural Models for Radiology Report Retrieval with Lexicon-based Automated Annotation
<https://aclanthology.org/2022.naacl-main.253>
AUTHORS: Luyao Shi, Tanveer Syeda-mahmood, Tyler Baldwin
HIGHLIGHT: In this paper we combine clinical finding detection with supervised query match learning.
- 254, TITLE: Transparent Human Evaluation for Image Captioning
<https://aclanthology.org/2022.naacl-main.254>
AUTHORS: Jungo Kasai, Keisuke Sakaguchi, Lavinia Dunagan, Jacob Morrison, Ronan Le Bras, Yejin Choi, Noah Smith
HIGHLIGHT: We establish THumB, a rubric-based human evaluation protocol for image captioning models.
- 255, TITLE: Lifting the Curse of Multilinguality by Pre-training Modular Transformers
<https://aclanthology.org/2022.naacl-main.255>
AUTHORS: Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, Mikel Artetxe
HIGHLIGHT: Multilingual pre-trained models are known to suffer from the curse of multilinguality, which causes per-language performance to drop as they cover more languages. We address this issue by introducing language-specific modules, which allows us to grow the total capacity of the model, while keeping the total number of trainable parameters per language constant.
- 256, TITLE: DocAMR: Multi-Sentence AMR Representation and Evaluation
<https://aclanthology.org/2022.naacl-main.256>
AUTHORS: Tahira Naseem, Austin Blodgett, Sadhana Kumaravel, Tim O’Gorman, Young-Suk Lee, Jeffrey Flanigan, Ram?n Astudillo, Radu Florian, Salim Roukos, Nathan Schneider
HIGHLIGHT: Taking advantage of a super-sentential level of coreference annotation from previous work, we introduce a simple algorithm for deriving a unified graph representation, avoiding the pitfalls of information loss from over-merging and lack of coherence from under merging.
- 257, TITLE: Learning to Transfer Prompts for Text Generation
<https://aclanthology.org/2022.naacl-main.257>
AUTHORS: Junyi Li, Tianyi Tang, Jian-Yun Nie, Ji-Rong Wen, Xin Zhao
HIGHLIGHT: In this paper, we improve this technique and propose a novel prompt-based method (PTG) for text generation in a transferable setting.
- 258, TITLE: ElitePLM: An Empirical Study on General Language Ability Evaluation of Pretrained Language Models
<https://aclanthology.org/2022.naacl-main.258>
AUTHORS: Junyi Li, Tianyi Tang, Zheng Gong, Lixin Yang, Zhuohao Yu, Zhipeng Chen, Jingyuan Wang, Xin Zhao, Ji-Rong Wen
HIGHLIGHT: In this paper, we present a large-scale empirical study on general language ability evaluation of PLMs (ElitePLM).
- 259, TITLE: Bidimensional Leaderboards: Generate and Evaluate Language Hand in Hand
<https://aclanthology.org/2022.naacl-main.259>
AUTHORS: Jungo Kasai, Keisuke Sakaguchi, Ronan Le Bras, Lavinia Dunagan, Jacob Morrison, Alexander Fabbri, Yejin Choi, Noah Smith
HIGHLIGHT: We argue that new advances on models and metrics should each more directly benefit and inform the other. We therefore propose a generalization of leaderboards, bidimensional leaderboards (Billboards), that simultaneously tracks progress in language generation models and metrics for their evaluation.
- 260, TITLE: Improving In-Context Few-Shot Learning via Self-Supervised Training
<https://aclanthology.org/2022.naacl-main.260>
AUTHORS: Mingda Chen, Jingfei Du, Ramakanth Pasunuru, Todor Mihaylov, Srini Iyer, Veselin Stoyanov, Zornitsa Kozareva
HIGHLIGHT: In this paper, we propose to use self-supervision in an intermediate training stage between pretraining and downstream few-shot usage with the goal to teach the model to perform in-context few shot learning.

- 261, TITLE: Exposing the Limits of Video-Text Models through Contrast Sets
<https://aclanthology.org/2022.naacl-main.261>
AUTHORS: Jae Sung Park, Sheng Shen, Ali Farhadi, Trevor Darrell, Yejin Choi, Anna Rohrbach
HIGHLIGHT: Can they discriminate between similar entities and actions? To answer this, we propose an evaluation framework that probes video-text models with hard negatives.
- 262, TITLE: Zero-shot Sonnet Generation with Discourse-level Planning and Aesthetics Features
<https://aclanthology.org/2022.naacl-main.262>
AUTHORS: Yufei Tian, Nanyun Peng
HIGHLIGHT: In this paper, we present a novel framework to generate sonnets that does not require training on poems.
- 263, TITLE: Benchmarking Intersectional Biases in NLP
<https://aclanthology.org/2022.naacl-main.263>
AUTHORS: John Lalor, Yi Yang, Kendall Smith, Nicole Forsgren, Ahmed Abbasi
HIGHLIGHT: In this work, we benchmark multiple NLP models with regards to their fairness and predictive performance across a variety of NLP tasks.
- 264, TITLE: When is BERT Multilingual? Isolating Crucial Ingredients for Cross-lingual Transfer
<https://aclanthology.org/2022.naacl-main.264>
AUTHORS: Ameet Deshpande, Partha Talukdar, Karthik Narasimhan
HIGHLIGHT: In this paper, we perform a large-scale empirical study to isolate the effects of various linguistic properties by measuring zero-shot transfer between four diverse natural languages and their counterparts constructed by modifying aspects such as the script, word order, and syntax.
- 265, TITLE: How Conservative are Language Models? Adapting to the Introduction of Gender-Neutral Pronouns
<https://aclanthology.org/2022.naacl-main.265>
AUTHORS: Stephanie Brandl, Ruixiang Cui, Anders Søgaard
HIGHLIGHT: We show that gender-neutral pronouns in Danish, English, and Swedish are associated with higher perplexity, more dispersed attention patterns, and worse downstream performance.
- 266, TITLE: Prompt Waywardness: The Curious Case of Discretized Interpretation of Continuous Prompts
<https://aclanthology.org/2022.naacl-main.266>
AUTHORS: Daniel Khashabi, Xinxu Lyu, Sewon Min, Lianhui Qin, Kyle Richardson, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, Sameer Singh, Yejin Choi
HIGHLIGHT: Motivated by these promising results, we investigate the feasibility of extracting a discrete (textual) interpretation of continuous prompts that is faithful to the problem they solve.
- 267, TITLE: Contrastive Representation Learning for Cross-Document Coreference Resolution of Events and Entities
<https://aclanthology.org/2022.naacl-main.267>
AUTHORS: Benjamin Hsu, Graham Horwood
HIGHLIGHT: We present an approach to entity and event coreference resolution utilizing contrastive representation learning.
- 268, TITLE: Learning the Ordering of Coordinate Compounds and Elaborate Expressions in Hmong, Lahu, and Chinese
<https://aclanthology.org/2022.naacl-main.268>
AUTHORS: Chenxuan Cui, Katherine Zhang, David Mortensen
HIGHLIGHT: We investigate whether the ordering of CCs and EEs can be learned empirically and whether computational models (classifiers and sequence-labeling models) learn unnatural hierarchies similar to those posited by Mortensen (2006).
- 269, TITLE: FRUIT: Faithfully Reflecting Updated Information in Text
<https://aclanthology.org/2022.naacl-main.269>
AUTHORS: Robert Iv, Alexandre Passos, Sameer Singh, Ming-Wei Chang
HIGHLIGHT: In this paper, we introduce the novel generation task of *faithfully reflecting updated information in text* (FRUIT) where the goal is to update an existing article given new evidence.
- 270, TITLE: Multi2WOZ: A Robust Multilingual Dataset and Conversational Pretraining for Task-Oriented Dialog
<https://aclanthology.org/2022.naacl-main.270>
AUTHORS: Chia-Chien Hung, Anne Lauscher, Ivan Vulic, Simone Ponzetto, Goran Glava?
HIGHLIGHT: In this work, we introduce Multi2WOZ, a new multilingual multi-domain TOD dataset, derived from the well-established English dataset MultiWOZ, that spans four typologically diverse languages: Chinese, German, Arabic, and Russian.
- 271, TITLE: ChapterBreak: A Challenge Dataset for Long-Range Language Models
<https://aclanthology.org/2022.naacl-main.271>
AUTHORS: Simeng Sun, Katherine Thai, Mohit Iyyer
HIGHLIGHT: To this end, we introduce ChapterBreak, a challenge dataset that provides an LRLM with a long segment from a narrative that ends at a chapter boundary and asks it to distinguish the beginning of the ground-truth next chapter from a set of negative segments sampled from the same narrative.
- 272, TITLE: CoBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction
<https://aclanthology.org/2022.naacl-main.272>

AUTHORS: Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, Matei Zaharia
HIGHLIGHT: In this work, we introduce Maize, a retriever that couples an aggressive residual compression mechanism with a denoised supervision strategy to simultaneously improve the quality and space footprint of late interaction.

273, **TITLE:** Quantifying Language Variation Acoustically with Few Resources
<https://aclanthology.org/2022.naacl-main.273>

AUTHORS: Martijn Bartelds, Martijn Wieling
HIGHLIGHT: However, deep acoustic models might have learned linguistic information that transfers to low-resource languages. In this study, we evaluate whether this is the case through the task of distinguishing low-resource (Dutch) regional varieties.

274, **TITLE:** Adaptable Adapters
<https://aclanthology.org/2022.naacl-main.274>

AUTHORS: Nafise Moosavi, Quentin Delfosse, Kristian Kersting, Iryna Gurevych
HIGHLIGHT: In this work, we introduce adaptable adapters that contain (1) learning different activation functions for different layers and different input data, and (2) a learnable switch to select and only use the beneficial adapter layers.

275, **TITLE:** Models in the Loop: Aiding Crowdworkers with Generative Annotation Assistants
<https://aclanthology.org/2022.naacl-main.275>

AUTHORS: Max Bartolo, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia, Douwe Kiela
HIGHLIGHT: In this work, we examine whether we can maintain the advantages of DADC, without incurring the additional cost.

276, **TITLE:** GMN: Generative Multi-modal Network for Practical Document Information Extraction
<https://aclanthology.org/2022.naacl-main.276>

AUTHORS: Haoyu Cao, Jiefeng Ma, Antai Guo, Yiqing Hu, Hao Liu, Deqiang Jiang, Yinsong Liu, Bo Ren
HIGHLIGHT: Although recent literature has already achieved competitive results, these approaches usually fail when dealing with complex documents with noisy OCR results or mutative layouts. This paper proposes Generative Multi-modal Network (GMN) for real-world scenarios to address these problems, which is a robust multi-modal generation method without predefined label categories.

277, **TITLE:** One Reference Is Not Enough: Diverse Distillation with Reference Selection for Non-Autoregressive Translation
<https://aclanthology.org/2022.naacl-main.277>

AUTHORS: Chenze Shao, Xuanfu Wu, Yang Feng
HIGHLIGHT: In this paper, we argue that one reference is not enough and propose diverse distillation with reference selection (DDRS) for NAT.

278, **TITLE:** Can Rationalization Improve Robustness?
<https://aclanthology.org/2022.naacl-main.278>

AUTHORS: Howard Chen, Jacqueline He, Karthik Narasimhan, Danqi Chen
HIGHLIGHT: A growing line of work has investigated the development of neural NLP models that can produce rationales-subsets of input that can explain their model predictions. In this paper, we ask whether such rationale models can provide robustness to adversarial attacks in addition to their interpretable nature.

279, **TITLE:** On the Effectiveness of Sentence Encoding for Intent Detection Meta-Learning
<https://aclanthology.org/2022.naacl-main.279>

AUTHORS: Tingting Ma, Qianhui Wu, Zhiwei Yu, Tiejun Zhao, Chin-Yew Lin
HIGHLIGHT: In this paper, we conduct empirical studies on a number of general-purpose sentence embedding schemes, showing that good sentence embeddings without any fine-tuning on intent detection data could produce a non-trivially strong performance.

280, **TITLE:** A Computational Acquisition Model for Multimodal Word Categorization
<https://aclanthology.org/2022.naacl-main.280>

AUTHORS: Uri Berger, Gabriel Stanovsky, Omri Abend, Lea Frermann
HIGHLIGHT: This is (a) not faithful to the information children receive and (b) prohibits the evaluation of such models with respect to category learning tasks, due to the pre-imposed category structure. We address this gap, and present a cognitively-inspired, multimodal acquisition model, trained from image-caption pairs on naturalistic data using cross-modal self-supervision.

281, **TITLE:** Residue-Based Natural Language Adversarial Attack Detection
<https://aclanthology.org/2022.naacl-main.281>

AUTHORS: Vyas Raina, Mark Gales
HIGHLIGHT: As an equivalent model-focused NLP detection approach, this work proposes a simple sentence-embedding "residue" based detector to identify adversarial examples.

282, **TITLE:** Does it Really Generalize Well on Unseen Data? Systematic Evaluation of Relational Triple Extraction Methods
<https://aclanthology.org/2022.naacl-main.282>

AUTHORS: Juhyuk Lee, Min-Joong Lee, June Yong Yang, Eunho Yang

HIGHLIGHT: To keep a knowledge graph up-to-date, an extractor needs not only the ability to recall the triples it encountered during training, but also the ability to extract the new triples from the context that it has never seen before. In this paper, we show that although existing extraction models are able to easily memorize and recall already seen triples, they cannot generalize effectively for unseen triples.

283, **TITLE:** From spoken dialogue to formal summary: An utterance rewriting for dialogue summarization
<https://aclanthology.org/2022.naacl-main.283>

AUTHORS: Yue Fang, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Bo Long, Yanyan Lan, Yanquan Zhou
HIGHLIGHT: However, the current state-of-the-art models pay more attention on the topic or structure of summary, rather than the consistency of dialogue summary with its input dialogue context, which may suffer from the personal and logical inconsistency problem. In this paper, we propose a new model, named ReWriteSum, to tackle this problem.

284, **TITLE:** EASE: Entity-Aware Contrastive Learning of Sentence Embedding
<https://aclanthology.org/2022.naacl-main.284>

AUTHORS: Sosuke Nishikawa, Ryokan Ri, Ikuya Yamada, Yoshimasa Tsuruoka, Isao Echizen
HIGHLIGHT: We present EASE, a novel method for learning sentence embeddings via contrastive learning between sentences and their related entities.

285, **TITLE:** Is Neural Topic Modelling Better than Clustering? An Empirical Study on Clustering with Contextual Embeddings for Topics
<https://aclanthology.org/2022.naacl-main.285>

AUTHORS: Zihan Zhang, Meng Fang, Ling Chen, Mohammad Reza Namazi Rad
HIGHLIGHT: In this paper, we conduct thorough experiments showing that directly clustering high-quality sentence embeddings with an appropriate word selecting method can generate more coherent and diverse topics than NTMs, achieving also higher efficiency and simplicity.

286, **TITLE:** Dynamic Multistep Reasoning based on Video Scene Graph for Video Question Answering
<https://aclanthology.org/2022.naacl-main.286>

AUTHORS: Jianguo Mao, Wenbin Jiang, Xiangdong Wang, Zhifan Feng, Yajuan Lyu, Hong Liu, Yong Zhu
HIGHLIGHT: We propose for video QA a novel model which performs dynamic multistep reasoning between questions and videos.

287, **TITLE:** TRUE: Re-evaluating Factual Consistency Evaluation
<https://aclanthology.org/2022.naacl-main.287>

AUTHORS: Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, Yossi Matias
HIGHLIGHT: In this work, we introduce TRUE: a comprehensive survey and assessment of factual consistency metrics on a standardized collection of existing texts from diverse tasks, manually annotated for factual consistency.

288, **TITLE:** Knowledge Inheritance for Pre-trained Language Models
<https://aclanthology.org/2022.naacl-main.288>

AUTHORS: Yujia Qin, Yankai Lin, Jing Yi, Jiajie Zhang, Xu Han, Zhengyan Zhang, Yusheng Su, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou
HIGHLIGHT: Specifically, we introduce a pre-training framework named "knowledge inheritance" (KI) and explore how could knowledge distillation serve as auxiliary supervision during pre-training to efficiently learn larger PLMs.

289, **TITLE:** Bi-SimCut: A Simple Strategy for Boosting Neural Machine Translation
<https://aclanthology.org/2022.naacl-main.289>

AUTHORS: Pengzhi Gao, Zhongjun He, Hua Wu, Haifeng Wang
HIGHLIGHT: We introduce Bi-SimCut: a simple but effective training strategy to boost neural machine translation (NMT) performance.

290, **TITLE:** On Transferability of Prompt Tuning for Natural Language Processing
<https://aclanthology.org/2022.naacl-main.290>

AUTHORS: Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun, Jie Zhou
HIGHLIGHT: To explore whether we can improve PT via prompt transfer, we empirically investigate the transferability of soft prompts across different downstream tasks and PLMs in this work.

291, **TITLE:** DocEE: A Large-Scale and Fine-grained Benchmark for Document-level Event Extraction
<https://aclanthology.org/2022.naacl-main.291>

AUTHORS: MeiHan Tong, Bin Xu, Shuai Wang, Meihuan Han, Yixin Cao, Jiangqi Zhu, Siyu Chen, Lei Hou, Juanzi Li
HIGHLIGHT: In this paper, we present DocEE, a new document-level event extraction dataset including 27,000+ events, 180,000+ arguments.

292, **TITLE:** Towards Debiasing Translation Artifacts
<https://aclanthology.org/2022.naacl-main.292>

AUTHORS: Koel Dutta Chowdhury, Rricha Jalota, Cristina Espa?a-Bonet, Josef Genabith
HIGHLIGHT: In this work, we propose a novel approach to reducing translationese by extending an established bias-removal technique.

- 293, TITLE: WECHSEL: Effective initialization of subword embeddings for cross-lingual transfer of monolingual language models
<https://aclanthology.org/2022.naacl-main.293>
AUTHORS: Benjamin Minixhofer, Fabian Paischer, Navid Rekabsaz
HIGHLIGHT: It is exceedingly expensive to train these models in other languages. To alleviate this problem, we introduce a novel method - called WECHSEL - to efficiently and effectively transfer pretrained LMs to new languages.
- 294, TITLE: A New Concept of Knowledge based Question Answering (KBQA) System for Multi-hop Reasoning
<https://aclanthology.org/2022.naacl-main.294>
AUTHORS: Yu Wang, V.srinivasan@samsung.com V.srinivasan@samsung.com, Hongxia Jin
HIGHLIGHT: In this paper, we introduce a new concept of KBQA system which can leverage multiple reasoning paths' information and only requires labeled answer as supervision.
- 295, TITLE: Bilingual Tabular Inference: A Case Study on Indic Languages
<https://aclanthology.org/2022.naacl-main.295>
AUTHORS: Chaitanya Agarwal, Vivek Gupta, Anoop Kunchukuttan, Manish Shrivastava
HIGHLIGHT: As a result, we present the challenging task of bilingual Tabular Natural Language Inference (bTNLI), in which the tabular premise and a hypothesis over it are in two separate languages.
- 296, TITLE: Generative Biomedical Entity Linking via Knowledge Base-Guided Pre-training and Synonyms-Aware Fine-tuning
<https://aclanthology.org/2022.naacl-main.296>
AUTHORS: Hongyi Yuan, Zheng Yuan, Sheng Yu
HIGHLIGHT: In this work, we use a generative approach to model biomedical EL and propose to inject synonyms knowledge in it.
- 297, TITLE: Robust Self-Augmentation for Named Entity Recognition with Meta Reweighting
<https://aclanthology.org/2022.naacl-main.297>
AUTHORS: Linzhi Wu, Pengjun Xie, Jie Zhou, Meishan Zhang, Ma Chunping, Guangwei Xu, Min Zhang
HIGHLIGHT: Prior research has mainly resorted to heuristic rule-based constraints to reduce the noise for specific self-augmentation methods individually. In this paper, we revisit these two typical self-augmentation methods for NER, and propose a unified meta-reweighting strategy for them to achieve a natural integration.
- 298, TITLE: Unsupervised Stem-based Cross-lingual Part-of-Speech Tagging for Morphologically Rich Low-Resource Languages
<https://aclanthology.org/2022.naacl-main.298>
AUTHORS: Ramy Eskander, Cass Lowry, Sujay Khandagale, Judith Klavans, Maria Polinsky, Smaranda Muresan
HIGHLIGHT: Our contributions are: 1) we propose an unsupervised stem-based cross-lingual approach for POS tagging for low-resource languages of rich morphology; 2) we further investigate morpheme-level alignment and projection; and 3) we examine whether the use of linguistic priors for morphological segmentation improves POS tagging.
- 299, TITLE: Optimising Equal Opportunity Fairness in Model Training
<https://aclanthology.org/2022.naacl-main.299>
AUTHORS: Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, Lea Frermann
HIGHLIGHT: In this work, we propose two novel training objectives which directly optimise for the widely-used criterion of equal opportunity, and show that they are effective in reducing bias while maintaining high performance over two classification tasks.
- 300, TITLE: Leaner and Faster: Two-Stage Model Compression for Lightweight Text-Image Retrieval
<https://aclanthology.org/2022.naacl-main.300>
AUTHORS: Siyu Ren, Kenny Zhu
HIGHLIGHT: In this paper, we present an effective two-stage framework to compress large pre-trained dual-encoder for lightweight text-image retrieval.
- 301, TITLE: Joint Learning-based Heterogeneous Graph Attention Network for Timeline Summarization
<https://aclanthology.org/2022.naacl-main.301>
AUTHORS: Jingyi You, Dongyuan Li, Hidetaka Kamigaito, Kotaro Funakoshi, Manabu Okumura
HIGHLIGHT: In this paper, we present a joint learning-based heterogeneous graph attention network for TLS (HeterTIs), in which date selection and event detection are combined into a unified framework to improve the extraction accuracy and remove redundant sentences simultaneously.
- 302, TITLE: Early Rumor Detection Using Neural Hawkes Process with a New Benchmark Dataset
<https://aclanthology.org/2022.naacl-main.302>
AUTHORS: Fengzhu Zeng, Wei Gao
HIGHLIGHT: Little attention has been paid on EARly Rumor Detection (EARD), and EARD performance was evaluated inappropriately on a few datasets where the actual early-stage information is largely missing. To reverse such situation, we construct BEARD, a new Benchmark dataset for EARD, based on claims from fact-checking websites by trying to gather as many early relevant posts as possible.
- 303, TITLE: Emp-RFT: Empathetic Response Generation via Recognizing Feature Transitions between Utterances

<https://aclanthology.org/2022.naacl-main.303>

AUTHORS: Wongyu Kim, Youbin Ahn, Donghyun Kim, Kyong-Ho Lee

HIGHLIGHT: However, existing approaches fail to perceive the transitions because they extract features for the context at the coarse-grained level. To solve the above issue, we propose a novel approach of recognizing feature transitions between utterances, which helps understand the dialogue flow and better grasp the features of utterance that needs attention.

304, TITLE: KCD: Knowledge Walks and Textual Cues Enhanced Political Perspective Detection in News Media

<https://aclanthology.org/2022.naacl-main.304>

AUTHORS: Wenqian Zhang, Shangbin Feng, Zilong Chen, Zhenyu Lei, Jundong Li, Minnan Luo

HIGHLIGHT: Previous approaches generally focus on leveraging textual content to identify stances, while they fail to reason with background knowledge or leverage the rich semantic and syntactic textual labels in news articles. In light of these limitations, we propose KCD, a political perspective detection approach to enable multi-hop knowledge reasoning and incorporate textual cues as paragraph-level labels.

305, TITLE: Collective Relevance Labeling for Passage Retrieval

<https://aclanthology.org/2022.naacl-main.305>

AUTHORS: Jihyuk Kim, Minsoo Kim, Seung-won Hwang

HIGHLIGHT: In contrast, we propose knowledge distillation for informed labeling, without incurring high computation overheads at evaluation time.

306, TITLE: COGMEN: COntextualized GNN based Multimodal Emotion recognitioN

<https://aclanthology.org/2022.naacl-main.306>

AUTHORS: Abhinav Joshi, Ashwani Bhat, Ayush Jain, Atin Singh, Ashutosh Modi

HIGHLIGHT: In this paper, we propose COntextualized Graph Neural Network based Multi-modal Emotion recognitioN (COGMEN) system that leverages local information (i.e., inter/intra dependency between speakers) and global information (context).

307, TITLE: Revisit Overconfidence for OOD Detection: Reassigned Contrastive Learning with Adaptive Class-dependent Threshold

<https://aclanthology.org/2022.naacl-main.307>

AUTHORS: Yanan Wu, Keqing He, Yuanmeng Yan, QiXiang Gao, Zhiyuan Zeng, Fujia Zheng, Lulu Zhao, Huixing Jiang, Wei Wu, Weiran Xu

HIGHLIGHT: In this paper, we comprehensively analyze overconfidence and classify it into two perspectives: over-confident OOD and in-domain (IND).

308, TITLE: AISFG: Abundant Information Slot Filling Generator

<https://aclanthology.org/2022.naacl-main.308>

AUTHORS: Yang Yan, Junda Ye, Zhongbao Zhang, Liwen Wang

HIGHLIGHT: Previous researches on zero/few-shot cross-domain slot filling focus on slot descriptions and examples while ignoring the slot type ambiguity and example ambiguity issues. To address these problems, we propose Abundant Information Slot Filling Generator (AISFG), a generative model with a novel query template that incorporates domain descriptions, slot descriptions, and examples with context.

309, TITLE: Improving negation detection with negation-focused pre-training

<https://aclanthology.org/2022.naacl-main.309>

AUTHORS: Thanh Truong, Timothy Baldwin, Trevor Cohn, Karin Verspoor

HIGHLIGHT: We propose a new negation-focused pre-training strategy, involving targeted data augmentation and negation masking, to better incorporate negation information into language models.

310, TITLE: Practice Makes a Solver Perfect: Data Augmentation for Math Word Problem Solvers

<https://aclanthology.org/2022.naacl-main.310>

AUTHORS: Vivek Kumar, Rishabh Maheshwary, Vikram Pudi

HIGHLIGHT: In this paper, we first conduct experiments to showcase that this behaviour is mainly associated with the limited size and diversity present in existing MWP datasets. Next, we propose several data augmentation techniques broadly categorized into Substitution and Paraphrasing based methods.

311, TITLE: DiffCSE: Difference-based Contrastive Learning for Sentence Embeddings

<https://aclanthology.org/2022.naacl-main.311>

AUTHORS: Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljagic, Shang-Wen Li, Scott Yih, Yoon Kim, James Glass

HIGHLIGHT: We propose DiffCSE, an unsupervised contrastive learning framework for learning sentence embeddings.

312, TITLE: Generative Cross-Domain Data Augmentation for Aspect and Opinion Co-Extraction

<https://aclanthology.org/2022.naacl-main.312>

AUTHORS: Junjie Li, Jianfei Yu, Rui Xia

HIGHLIGHT: In this paper, we propose a new Generative Cross-Domain Data Augmentation framework for unsupervised domain adaptation.

313, TITLE: ProQA: Structural Prompt-based Pre-training for Unified Question Answering

<https://aclanthology.org/2022.naacl-main.313>

AUTHORS: Wanjun Zhong, Yifan Gao, Ning Ding, Yujia Qin, Zhiyuan Liu, Ming Zhou, Jiahai Wang, Jian Yin, Nan Duan

- HIGHLIGHT:** The specialty in QA research hinders systems from modeling commonalities between tasks and generalization for wider applications. To address this issue, we present ProQA, a unified QA paradigm that solves various tasks through a single model.
- 314, **TITLE:** A Data Cartography based MixUp for Pre-trained Language Models
<https://aclanthology.org/2022.naacl-main.314>
AUTHORS: Seo Yeon Park, Cornelia Caragea
HIGHLIGHT: In this work, we propose TDMixUp, a novel MixUp strategy that leverages Training Dynamics and allows more informative samples to be combined for generating new data samples.
- 315, **TITLE:** Grapheme-to-Phoneme Conversion for Thai using Neural Regression Models
<https://aclanthology.org/2022.naacl-main.315>
AUTHORS: Tomohiro Yamasaki
HIGHLIGHT: We propose a novel Thai grapheme-to-phoneme conversion method based on a neural regression model that is trained using neural networks to predict the similarity between a candidate and the correct pronunciation.
- 316, **TITLE:** Generating Authentic Adversarial Examples beyond Meaning-preserving with Doubly Round-trip Translation
<https://aclanthology.org/2022.naacl-main.316>
AUTHORS: Siyu Lai, Zhen Yang, Fandong Meng, Xue Zhang, Yufeng Chen, Jinan Xu, Jie Zhou
HIGHLIGHT: However, a potential pitfall for this approach is that we cannot decide whether the generated examples are adversarial to the target NMT model or the auxiliary backward one, as the reconstruction error through the RTT can be related to either. To remedy this problem, we propose a new definition for NMT adversarial examples based on the Doubly Round-Trip Translation (DRTT).
- 317, **TITLE:** TVShowGuess: Character Comprehension in Stories as Speaker Guessing
<https://aclanthology.org/2022.naacl-main.317>
AUTHORS: Yisi Sang, Xiangyang Mou, Mo Yu, Shunyu Yao, Jing Li, Jeffrey Stanton
HIGHLIGHT: We propose a new task for assessing machines' skills of understanding fictional characters in narrative stories.
- 318, **TITLE:** Causal Distillation for Language Models
<https://aclanthology.org/2022.naacl-main.318>
AUTHORS: Zhengxuan Wu, Atticus Geiger, Joshua Rozner, Elisa Kreiss, Hanson Lu, Thomas Icard, Christopher Potts, Noah Goodman
HIGHLIGHT: In this paper, we show that it is beneficial to augment distillation with a third objective that encourages the student to imitate the causal dynamics of the teacher through a distillation interchange intervention training objective (DIITO).
- 319, **TITLE:** FNet: Mixing Tokens with Fourier Transforms
<https://aclanthology.org/2022.naacl-main.319>
AUTHORS: James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, Santiago Ontanon
HIGHLIGHT: We show that Transformer encoder architectures can be sped up, with limited accuracy costs, by replacing the self-attention sublayers with simple linear transformations that "mix" input tokens.
- 320, **TITLE:** Answer Consolidation: Formulation and Benchmarking
<https://aclanthology.org/2022.naacl-main.320>
AUTHORS: Wenxuan Zhou, Qiang Ning, Heba Elfardy, Kevin Small, Muhao Chen
HIGHLIGHT: In this paper, we formulate the problem of answer consolidation, where answers are partitioned into multiple groups, each representing different aspects of the answer set.
- 321, **TITLE:** Informativeness and Invariance: Two Perspectives on Spurious Correlations in Natural Language
<https://aclanthology.org/2022.naacl-main.321>
AUTHORS: Jacob Eisenstein
HIGHLIGHT: Gardner et al (2021) argue that the compositional nature of language implies that all correlations between labels and individual input features are spurious. This paper analyzes this proposal in the context of a toy example, demonstrating three distinct conditions that can give rise to feature-label correlations in a simple PCFG.
- 322, **TITLE:** FOAM: A Follower-aware Speaker Model For Vision-and-Language Navigation
<https://aclanthology.org/2022.naacl-main.322>
AUTHORS: Zi-Yi Dou, Nanyun Peng
HIGHLIGHT: In this paper, we present FOAM, a Follower-Aware speaker Model that is constantly updated given the follower feedback, so that the generated instructions can be more suitable to the current learning state of the follower.
- 323, **TITLE:** Improving Compositional Generalization with Latent Structure and Data Augmentation
<https://aclanthology.org/2022.naacl-main.323>
AUTHORS: Linlu Qiu, Peter Shaw, Panupong Pasupat, Pawel Nowak, Tal Linzen, Fei Sha, Kristina Toutanova
HIGHLIGHT: We present a more powerful data recombination method using a model called Compositional Structure Learner (CSL).
- 324, **TITLE:** Joint Extraction of Entities, Relations, and Events via Modeling Inter-Instance and Inter-Label Dependencies
<https://aclanthology.org/2022.naacl-main.324>
AUTHORS: Minh Van Nguyen, Bonan Min, Franck Dernoncourt, Thien Nguyen

HIGHLIGHT: However, previous JointIE models often assume heuristic manually-designed dependency between the task instances and mean-field factorization for the joint distribution of instance labels, thus unable to capture optimal dependencies among instances and labels to improve representation learning and IE performance. To overcome these limitations, we propose to induce a dependency graph among task instances from data to boost representation learning.

325, **TITLE:** Linguistic Frameworks Go Toe-to-Toe at Neuro-Symbolic Language Modeling
<https://aclanthology.org/2022.naacl-main.325>

AUTHORS: Jakob Prange, Nathan Schneider, Lingpeng Kong

HIGHLIGHT: We examine the extent to which, in principle, different syntactic and semantic graph representations can complement and improve neural language modeling.

326, **TITLE:** Imagination-Augmented Natural Language Understanding
<https://aclanthology.org/2022.naacl-main.326>

AUTHORS: Yujie Lu, Wanrong Zhu, Xin Wang, Miguel Eckstein, William Yang Wang

HIGHLIGHT: Therefore, we introduce an Imagination-Augmented Cross-modal Encoder (iACE) to solve natural language understanding tasks from a novel learning perspective-imagination-augmented cross-modal understanding.

327, **TITLE:** What company do words keep? Revisiting the distributional semantics of J.R. Firth & Zellig Harris

<https://aclanthology.org/2022.naacl-main.327>

AUTHORS: Mikael Brunila, Jack LaViolette

HIGHLIGHT: Contrasting these theories from the perspective of current debates in NLP, we discover in Firth a figure who could guide the field towards a more culturally grounded notion of semantics.

328, **TITLE:** Compositional Task-Oriented Parsing as Abstractive Question Answering

<https://aclanthology.org/2022.naacl-main.328>

AUTHORS: Wenting Zhao, Konstantine Arkoudas, Weiqi Sun, Claire Cardie

HIGHLIGHT: In this work we continue to explore naturalized semantic parsing by presenting a general reduction of TOP to abstractive question answering that overcomes some limitations of canonical paraphrasing.

329, **TITLE:** Learning Cross-Lingual IR from an English Retriever

<https://aclanthology.org/2022.naacl-main.329>

AUTHORS: Yulong Li, Martin Franz, Md Arafat Sultan, Bhavani Iyer, Young-Suk Lee, Avirup Sil

HIGHLIGHT: We present DR.DECR (Dense Retrieval with Distillation-Enhanced Cross-Lingual Representation), a new cross-lingual information retrieval (CLIR) system trained using multi-stage knowledge distillation (KD).

330, **TITLE:** Testing the Ability of Language Models to Interpret Figurative Language

<https://aclanthology.org/2022.naacl-main.330>

AUTHORS: Emmy Liu, Chenxuan Cui, Kenneth Zheng, Graham Neubig

HIGHLIGHT: However, figurative language has been a relatively under-studied area in NLP, and it remains an open question to what extent modern language models can interpret nonliteral phrases. To address this question, we introduce Fig-QA, a Winograd-style nonliteral language understanding task consisting of correctly interpreting paired figurative phrases with divergent meanings.

331, **TITLE:** Multi-Vector Models with Textual Guidance for Fine-Grained Scientific Document Similarity

<https://aclanthology.org/2022.naacl-main.331>

AUTHORS: Sheshera Mysore, Arman Cohan, Tom Hope

HIGHLIGHT: We present a new scientific document similarity model based on matching fine-grained aspects of texts.

332, **TITLE:** CHAI: A CHatbot AI for Task-Oriented Dialogue with Offline Reinforcement Learning

<https://aclanthology.org/2022.naacl-main.332>

AUTHORS: Siddharth Verma, Justin Fu, Sherry Yang, Sergey Levine

HIGHLIGHT: In this paper, we study how offline reinforcement learning can instead be used to train dialogue agents entirely using static datasets collected from human speakers.

333, **TITLE:** Connecting the Dots between Audio and Text without Parallel Data through Visual Knowledge Transfer

<https://aclanthology.org/2022.naacl-main.333>

AUTHORS: Yanpeng Zhao, Jack Hessel, Youngjae Yu, Ximing Lu, Rowan Zellers, Yejin Choi

HIGHLIGHT: We propose VIP-ANT that induces Audio-Text alignment without using any parallel audio-text data.

334, **TITLE:** SURF: Semantic-level Unsupervised Reward Function for Machine Translation

<https://aclanthology.org/2022.naacl-main.334>

AUTHORS: Atijit Anuchitanukul, Julia Ive

HIGHLIGHT: This is due to the intrinsic difficulty of the task in the high-dimensional discrete action space as well as the sparseness of the standard reward functions defined for limited set of ground-truth sequences biased towards singular lexical choices. To address this issue, we formulate SURF, a maximally dense semantic-level unsupervised reward function which mimics human evaluation by considering both sentence fluency and semantic similarity.

335, **TITLE:** Disentangling Categorization in Multi-agent Emergent Communication

<https://aclanthology.org/2022.naacl-main.335>

AUTHORS: Washington Garcia, Hamilton Clouse, Kevin Butler

HIGHLIGHT: In this work, we propose the use of disentangled representations from representation learning to quantify the categorization power of agents, enabling a differential analysis between combinations of heterogeneous systems, e.g., pairs of agents which learn to communicate despite mismatched concept realization.

336, **TITLE:** Show, Don't Tell: Demonstrations Outperform Descriptions for Schema-Guided Task-Oriented Dialogue
<https://aclanthology.org/2022.naacl-main.336>

AUTHORS: Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi, Yonghui Wu

HIGHLIGHT: In this work, we propose Show, Don't Tell, which prompts seq2seq models with a labeled example dialogue to show the semantics of schema elements rather than tell the model through descriptions.

337, **TITLE:** Does Pre-training Induce Systematic Inference? How Masked Language Models Acquire Commonsense Knowledge
<https://aclanthology.org/2022.naacl-main.337>

AUTHORS: Ian Porada, Alessandro Sordani, Jackie Cheung

HIGHLIGHT: Transformer models pre-trained with a masked-language-modeling objective (e.g., BERT) encode commonsense knowledge as evidenced by behavioral probes; however, the extent to which this knowledge is acquired by systematic inference over the semantics of the pre-training corpora is an open question. To answer this question, we selectively inject verbalized knowledge into the pre-training minibatches of BERT and evaluate how well the model generalizes to supported inferences after pre-training on the injected knowledge.

338, **TITLE:** Using Paraphrases to Study Properties of Contextual Embeddings
<https://aclanthology.org/2022.naacl-main.338>

AUTHORS: Laura Burdick, Jonathan Kummerfeld, Rada Mihalcea

HIGHLIGHT: We use paraphrases as a unique source of data to analyze contextualized embeddings, with a particular focus on BERT.

HIGHLIGHT: We use paraphrases as a unique source of data to analyze contextualized embeddings, with a particular focus on BERT.

339, **TITLE:** Measure and Improve Robustness in NLP Models: A Survey
<https://aclanthology.org/2022.naacl-main.339>

AUTHORS: Xuezhi Wang, Haohan Wang, Diyi Yang

HIGHLIGHT: In this paper, we aim to provide a unifying survey of how to define, measure and improve robustness in NLP.

HIGHLIGHT: In this paper, we aim to provide a unifying survey of how to define, measure and improve robustness in NLP.

340, **TITLE:** Learning to Generate Examples for Semantic Processing Tasks
<https://aclanthology.org/2022.naacl-main.340>

AUTHORS: Danilo Croce, Simone Filice, Giuseppe Castellucci, Roberto Basili

HIGHLIGHT: In this paper, we propose a neural approach to automatically learn to generate new examples using a pre-trained sequence-to-sequence model.

HIGHLIGHT: In this paper, we propose a neural approach to automatically learn to generate new examples using a pre-trained sequence-to-sequence model.

341, **TITLE:** Symbolic Knowledge Distillation: from General Language Models to Commonsense Models
<https://aclanthology.org/2022.naacl-main.341>

AUTHORS: Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, Yejin Choi

HIGHLIGHT: In this work, we investigate an alternative, from-machine-to-corpus-to-machine: general language models author these commonsense knowledge graphs to train commonsense models.

HIGHLIGHT: In this work, we investigate an alternative, from-machine-to-corpus-to-machine: general language models author these commonsense knowledge graphs to train commonsense models.

HIGHLIGHT: In this work, we investigate an alternative, from-machine-to-corpus-to-machine: general language models author these commonsense knowledge graphs to train commonsense models.

342, **TITLE:** GenIE: Generative Information Extraction
<https://aclanthology.org/2022.naacl-main.342>

AUTHORS: Martin Josifoski, Nicola De Cao, Maxime Peyrard, Fabio Petroni, Robert West

HIGHLIGHT: We introduce GenIE (generative information extraction), the first end-to-end autoregressive formulation of closed information extraction.

HIGHLIGHT: We introduce GenIE (generative information extraction), the first end-to-end autoregressive formulation of closed information extraction.

343, **TITLE:** Entity Linking via Explicit Mention-Mention Coreference Modeling
<https://aclanthology.org/2022.naacl-main.343>

AUTHORS: Dhruv Agarwal, Rico Angell, Nicholas Monath, Andrew McCallum

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

HIGHLIGHT: In this paper, we present and empirically analyze a novel training approach for learning mention and entity representations that is based on building minimum spanning arborescences (i.e., directed spanning trees) over mentions and entities across documents to explicitly model mention coreference relationships.

- AUTHORS: Belinda Li, Jane Yu, Madian Khabsa, Luke Zettlemoyer, Alon Halevy, Jacob Andreas
HIGHLIGHT: We present a large-scale empirical study of the features and limits of LM adaptability using a new benchmark, TaskBench500, built from 500 procedurally generated sequence modeling tasks.
- 347, TITLE: Counterfactually Augmented Data and Unintended Bias: The Case of Sexism and Hate Speech Detection
<https://aclanthology.org/2022.naacl-main.347>
AUTHORS: Indira Sen, Mattia Samory, Claudia Wagner, Isabelle Augenstein
HIGHLIGHT: Here, we test models for sexism and hate speech detection on challenging data: non-hate and non-sexist usage of identity and gendered terms.
- 348, TITLE: A Study of the Attention Abnormality in Trojaned BERTs
<https://aclanthology.org/2022.naacl-main.348>
AUTHORS: Weimin Lyu, Songzhu Zheng, Tengfei Ma, Chao Chen
HIGHLIGHT: In this paper, we investigate the underlying mechanism of Trojaned BERT models.
- 349, TITLE: EPiDA: An Easy Plug-in Data Augmentation Framework for High Performance Text Classification
<https://aclanthology.org/2022.naacl-main.349>
AUTHORS: Minyi Zhao, Lu Zhang, Yi Xu, Jiandong Ding, Jihong Guan, Shuigeng Zhou
HIGHLIGHT: In this paper, we present an easy and plug-in data augmentation framework EPiDA to support effective text classification.
- 350, TITLE: Partial-input baselines show that NLI models can ignore context, but they don't.
<https://aclanthology.org/2022.naacl-main.350>
AUTHORS: Neha Srikanth, Rachel Rudinger
HIGHLIGHT: We introduce an evaluation set of 600 examples consisting of perturbed premises to examine a RoBERTa model's sensitivity to edited contexts.
- 351, TITLE: Lifelong Pretraining: Continually Adapting Language Models to Emerging Corpora
<https://aclanthology.org/2022.naacl-main.351>
AUTHORS: Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, Xiang Ren
HIGHLIGHT: In this paper, we study a lifelong language model pretraining challenge where a PTLM is continually updated so as to adapt to emerging data.
- 352, TITLE: Learning as Conversation: Dialogue Systems Reinforced for Information Acquisition
<https://aclanthology.org/2022.naacl-main.352>
AUTHORS: Pengshan Cai, Hui Wan, Fei Liu, Mo Yu, Hong Yu, Sachindra Joshi
HIGHLIGHT: We propose novel AI-empowered chat bots for learning as conversation where a user does not read a passage but gains information and knowledge through conversation with a teacher bot.
- 353, TITLE: Dynamic Programming in Rank Space: Scaling Structured Inference with Low-Rank HMMs and PCFGs
<https://aclanthology.org/2022.naacl-main.353>
AUTHORS: Songlin Yang, Wei Liu, Kewei Tu
HIGHLIGHT: However, inference with large state spaces is computationally demanding, especially for PCFGs. To tackle this challenge, we leverage tensor rank decomposition (aka. CPD) to decrease inference computational complexities for a subset of FGGs subsuming HMMs and PCFGs.
- 354, TITLE: What Factors Should Paper-Reviewer Assignments Rely On? Community Perspectives on Issues and Ideals in Conference Peer-Review
<https://aclanthology.org/2022.naacl-main.354>
AUTHORS: Terne Thorn Jakobsen, Anna Rogers
HIGHLIGHT: We present the results of the first survey of the NLP community, identifying common issues and perspectives on what factors should be considered by paper-reviewer matching systems.
- 355, TITLE: Reducing Disambiguation Biases in NMT by Leveraging Explicit Word Sense Information
<https://aclanthology.org/2022.naacl-main.355>
AUTHORS: Niccolò Campolungo, Tommaso Pasini, Denis Emelin, Roberto Navigli
HIGHLIGHT: In this paper, we first provide a novel approach for automatically creating high-precision sense-annotated parallel corpora, and then put forward a specifically tailored fine-tuning strategy for exploiting these sense annotations during training without introducing any additional requirement at inference time.
- 356, TITLE: Mining Clues from Incomplete Utterance: A Query-enhanced Network for Incomplete Utterance Rewriting
<https://aclanthology.org/2022.naacl-main.356>
AUTHORS: Shuzheng Si, Shuang Zeng, Baobao Chang
HIGHLIGHT: However, previous works do not consider the semantic structural information between incomplete utterance and rewritten utterance or model the semantic structure implicitly and insufficiently. To address this problem, we propose a QUERY-Enhanced Network(QUEEN) to solve this problem.
- 357, TITLE: Domain-Oriented Prefix-Tuning: Towards Efficient and Generalizable Fine-tuning for Zero-Shot Dialogue Summarization
<https://aclanthology.org/2022.naacl-main.357>

AUTHORS: Lulu Zhao, Fujia Zheng, Weihao Zeng, Keqing He, Weiran Xu, Huixing Jiang, Wei Wu, Yanan Wu
HIGHLIGHT: To explore the lightweight fine-tuning methods for domain adaptation of dialogue summarization, in this paper, we propose an efficient and generalizable Domain-Oriented Prefix-tuning model, which utilizes a domain word initialized prefix module to alleviate domain entanglement and adopts discrete prompts to guide the model to focus on key contents of dialogues and enhance model generalization.

358, **TITLE:** Interactive Symbol Grounding with Complex Referential Expressions
<https://aclanthology.org/2022.naacl-main.358>

AUTHORS: Rimvydas Rubavicius, Alex Lascarides
HIGHLIGHT: We present a procedure for learning to ground symbols from a sequence of stimuli consisting of an arbitrarily complex noun phrase (e.g. "all but one green square above both red circles.")

359, **TITLE:** Generalized Quantifiers as a Source of Error in Multilingual NLU Benchmarks
<https://aclanthology.org/2022.naacl-main.359>

AUTHORS: Ruixiang Cui, Daniel Hershcovich, Anders S?gaard
HIGHLIGHT: To facilitate directly-targeted probing, we present an adversarial generalized quantifier NLI task (GQNLI) and show that pre-trained language models have a clear lack of robustness in generalized quantifier reasoning.

360, **TITLE:** Exact Paired-Permutation Testing for Structured Test Statistics
<https://aclanthology.org/2022.naacl-main.360>

AUTHORS: Ran Zmigrod, Tim Vieira, Ryan Cotterell
HIGHLIGHT: In this paper, we provide an efficient exact algorithm for the paired-permutation test for a family of structured test statistics.

361, **TITLE:** A Balanced Data Approach for Evaluating Cross-Lingual Transfer: Mapping the Linguistic Blood Bank
<https://aclanthology.org/2022.naacl-main.361>

AUTHORS: Dan Malkin, Tomasz Limisiewicz, Gabriel Stanovsky
HIGHLIGHT: We inspect zero-shot performance in balanced data conditions to mitigate data size confounds, classifying pretraining languages that improve downstream performance as donors, and languages that are improved in zero-shot performance as recipients.

362, **TITLE:** SSEGCN: Syntactic and Semantic Enhanced Graph Convolutional Network for Aspect-based Sentiment Analysis
<https://aclanthology.org/2022.naacl-main.362>

AUTHORS: Zheng Zhang, Zili Zhou, Yanna Wang
HIGHLIGHT: In this paper, we propose a novel Syntactic and Semantic Enhanced Graph Convolutional Network (SSEGCN) model for ABSA task.

363, **TITLE:** Mitigating Toxic Degeneration with Empathetic Data: Exploring the Relationship Between Toxicity and Empathy
<https://aclanthology.org/2022.naacl-main.363>

AUTHORS: Allison Lahnala, Charles Welch, B?la Neuendorf, Lucie Flek
HIGHLIGHT: Using empathetic data, we improve over recent work on controllable text generation that aims to reduce the toxicity of generated text.

364, **TITLE:** DUCK: Rumour Detection on Social Media by Modelling User and Comment Propagation Networks
<https://aclanthology.org/2022.naacl-main.364>

AUTHORS: Lin Tian, Xiuzhen Zhang, Jey Han Lau
HIGHLIGHT: Motivated by the observation that the user network - which captures \textit{who} engage with a story - and the comment network - which captures \textit{how} they react to it - provide complementary signals for rumour detection, in this paper, we propose DUCK (rumour _detection with _user and _comment networ _ks) for rumour detection on social media.

365, **TITLE:** Jam or Cream First? Modeling Ambiguity in Neural Machine Translation with SCONES
<https://aclanthology.org/2022.naacl-main.365>

AUTHORS: Felix Stahlberg, Shankar Kumar
HIGHLIGHT: Machine translation, however, is intrinsically uncertain: the same source sentence can have multiple semantically equivalent translations. Therefore, we propose to replace the softmax activation with a multi-label classification layer that can model ambiguity more effectively.

366, **TITLE:** SkillSpan: Hard and Soft Skill Extraction from English Job Postings
<https://aclanthology.org/2022.naacl-main.366>

AUTHORS: Mike Zhang, Kristian Jensen, Sif Sonniks, Barbara Plank
HIGHLIGHT: To address this gap, we introduce SKILLSPAN, a novel SE dataset consisting of 14.5K sentences and over 12.5K annotated spans.

367, **TITLE:** RAAT: Relation-Augmented Attention Transformer for Relation Modeling in Document-Level Event Extraction
<https://aclanthology.org/2022.naacl-main.367>

AUTHORS: Yuan Liang, Zhuoxuan Jiang, Di Yin, Bo Ren

HIGHLIGHT: In this paper, we argue that the relation information of event arguments is of great significance for addressing the above two issues, and propose a new DEE framework which can model the relation dependencies, called Relation-augmented Document-level Event Extraction (ReDEE).

368, **TITLE:** A Double-Graph Based Framework for Frame Semantic Parsing
<https://aclanthology.org/2022.naacl-main.368>
AUTHORS: Ce Zheng, Xudong Chen, Runxin Xu, Baobao Chang
HIGHLIGHT: In this paper, we propose a Knowledge-guided Incremental semantic parser with Double-graph (KID).

369, **TITLE:** An Enhanced Span-based Decomposition Method for Few-Shot Sequence Labeling
<https://aclanthology.org/2022.naacl-main.369>
AUTHORS: Peiyi Wang, Runxin Xu, Tianyu Liu, Qingyu Zhou, Yunbo Cao, Baobao Chang, Zhifang Sui
HIGHLIGHT: However, most prior works assign a label to each token based on the token-level similarities, which ignores the integrality of named entities or slots. To this end, in this paper, we propose ESD, an Enhanced Span-based Decomposition method for FSSL.

370, **TITLE:** A Two-Stream AMR-enhanced Model for Document-level Event Argument Extraction
<https://aclanthology.org/2022.naacl-main.370>
AUTHORS: Runxin Xu, Peiyi Wang, Tianyu Liu, Shuang Zeng, Baobao Chang, Zhifang Sui
HIGHLIGHT: In this paper, we focus on extracting event arguments from an entire document, which mainly faces two critical problems: a) the long-distance dependency between trigger and arguments over sentences; b) the distracting context towards an event in the document.

371, **TITLE:** Robust (Controlled) Table-to-Text Generation with Structure-Aware Equivariance Learning
<https://aclanthology.org/2022.naacl-main.371>
AUTHORS: Fei Wang, Zhewei Xu, Pedro Szekely, Muhao Chen
HIGHLIGHT: Accordingly, we propose an equivariance learning framework, which encodes tables with a structure-aware self-attention mechanism.

372, **TITLE:** JointLK: Joint Reasoning with Language Models and Knowledge Graphs for Commonsense Question Answering
<https://aclanthology.org/2022.naacl-main.372>
AUTHORS: Yueqing Sun, Qi Shi, Le Qi, Yu Zhang
HIGHLIGHT: However, they ignore (i) the effectively fusing and reasoning over question context representations and the KG representations, and (ii) automatically selecting relevant nodes from the noisy KGs during reasoning. In this paper, we propose a novel model, JointLK, which solves the above limitations through the joint reasoning of LM and GNN and the dynamic KGs pruning mechanism.

373, **TITLE:** Models In a Spelling Bee: Language Models Implicitly Learn the Character Composition of Tokens
<https://aclanthology.org/2022.naacl-main.373>
AUTHORS: Itay Itzhak, Omer Levy
HIGHLIGHT: We probe the embedding layer of pretrained language models and show that models learn the internal character composition of whole word and subword tokens to a surprising extent, without ever seeing the characters coupled with the tokens.

374, **TITLE:** A Corpus for Understanding and Generating Moral Stories
<https://aclanthology.org/2022.naacl-main.374>
AUTHORS: Jian Guan, Ziqi Liu, Minlie Huang
HIGHLIGHT: Its challenges mainly lie in: (1) grasping knowledge about abstract concepts in morals, (2) capturing inter-event discourse relations in stories, and (3) aligning value preferences of stories and morals concerning good or bad behavior. In this paper, we propose two understanding tasks and two generation tasks to assess these abilities of machines.

375, **TITLE:** Modeling Multi-Granularity Hierarchical Features for Relation Extraction
<https://aclanthology.org/2022.naacl-main.375>
AUTHORS: Xinnian Liang, Shuangzhi Wu, Mu Li, Zhoujun Li
HIGHLIGHT: In this paper, we propose a novel method to extract multi-granularity features based solely on the original input sentences.

376, **TITLE:** Cross-modal Contrastive Learning for Speech Translation
<https://aclanthology.org/2022.naacl-main.376>
AUTHORS: Rong Ye, Mingxuan Wang, Lei Li
HIGHLIGHT: Learning similar representations for semantically similar speech and text is important for speech translation. To this end, we propose ConST, a cross-modal contrastive learning method for end-to-end speech-to-text translation.

377, **TITLE:** Meet Your Favorite Character: Open-domain Chatbot Mimicking Fictional Characters with only a Few Utterances
<https://aclanthology.org/2022.naacl-main.377>
AUTHORS: Seungju Han, Beomsu Kim, Jin Yong Yoo, Seokjun Seo, Sangbum Kim, Enkhbayar Erdence, Buru Chang
HIGHLIGHT: In this paper, we consider mimicking fictional characters as a promising direction for building engaging conversation models.

- 378, TITLE: DynamicTOC: Persona-based Table of Contents for Consumption of Long Documents
<https://aclanthology.org/2022.naacl-main.378>
AUTHORS: Himanshu Maheshwari, Nethraa Sivakumar, Shelly Jain, Tanvi Karandikar, Vinay Aggarwal, Navita Goyal, Sumit Shekhar
HIGHLIGHT: In this work, we describe DynamicToC, a dynamic table of content-based navigator, to aid in the task of non-linear, persona-based document consumption.
- 379, TITLE: KALA: Knowledge-Augmented Language Model Adaptation
<https://aclanthology.org/2022.naacl-main.379>
AUTHORS: Minki Kang, Jinheon Baek, Sung Ju Hwang
HIGHLIGHT: Moreover, adaptive pre-training can harm the PLM's performance on the downstream task by causing catastrophic forgetting of its general knowledge. To overcome such limitations of adaptive pre-training for PLM adaption, we propose a novel domain adaption framework for PLMs coined as Knowledge-Augmented Language model Adaptation (KALA), which modulates the intermediate hidden representations of PLMs with domain knowledge, consisting of entities and their relational facts.
- 380, TITLE: On the Effect of Pretraining Corpora on In-context Learning by a Large-scale Language Model
<https://aclanthology.org/2022.naacl-main.380>
AUTHORS: Seongjin Shin, Sang-Woo Lee, Hwijee Ahn, Sungdong Kim, HyoungSeok Kim, Boseop Kim, Kyunghyun Cho, Gichang Lee, Woomyoung Park, Jung-Woo Ha, Nako Sung
HIGHLIGHT: Here, we investigate the effects of the source and size of the pretraining corpus on in-context learning in HyperCLOVA, a Korean-centric GPT-3 model.
- 381, TITLE: Sketching as a Tool for Understanding and Accelerating Self-attention for Long Sequences
<https://aclanthology.org/2022.naacl-main.381>
AUTHORS: Yifan Chen, Qi Zeng, Dilek Hakkani-Tur, Di Jin, Heng Ji, Yun Yang
HIGHLIGHT: To address this limitation, Linformer and Informer reduce the quadratic complexity to linear (modulo logarithmic factors) via low-dimensional projection and row selection, respectively. These two models are intrinsically connected, and to understand their connection we introduce a theoretical framework of matrix sketching.
- 382, TITLE: Partner Personas Generation for Dialogue Response Generation
<https://aclanthology.org/2022.naacl-main.382>
AUTHORS: Hongyuan Lu, Wai Lam, Hong Cheng, Helen Meng
HIGHLIGHT: Moreover, in practical applications, the availability of the gold partner personas is often not the case. This paper attempts to tackle these issues by offering a novel framework that leverages automatic partner personas generation to enhance the succeeding dialogue response generation.
- 383, TITLE: Semantically Informed Slang Interpretation
<https://aclanthology.org/2022.naacl-main.383>
AUTHORS: Zhewei Sun, Richard Zemel, Yang Xu
HIGHLIGHT: We propose a semantically informed slang interpretation (SSI) framework that considers jointly the contextual and semantic appropriateness of a candidate interpretation for a query slang.
- 384, TITLE: Dual-Channel Evidence Fusion for Fact Verification over Texts and Tables
<https://aclanthology.org/2022.naacl-main.384>
AUTHORS: Nan Hu, Zirui Wu, Yuxuan Lai, Xiao Liu, Yansong Feng
HIGHLIGHT: In this paper, we propose a Dual Channel Unified Format fact verification model (DCUF), which unifies various evidence into parallel streams, i.e., natural language sentences and a global evidence table, simultaneously.
- 385, TITLE: TreeMix: Compositional Constituency-based Data Augmentation for Natural Language Understanding
<https://aclanthology.org/2022.naacl-main.385>
AUTHORS: Le Zhang, Zichao Yang, Diyi Yang
HIGHLIGHT: Though effective, they missed one important characteristic of language compositionality, meaning of a complex expression is built from its sub-parts. Motivated by this, we propose a compositional data augmentation approach for natural language understanding called TreeMix.
- 386, TITLE: Syn2Vec: Synset Colexification Graphs for Lexical Semantic Similarity
<https://aclanthology.org/2022.naacl-main.386>
AUTHORS: John Harvill, Roxana Girju, Mark Hasegawa-Johnson
HIGHLIGHT: In this paper we focus on patterns of colexification (co-expressions of form-meaning mapping in the lexicon) as an aspect of lexical-semantic organization, and use them to build large scale synset graphs across BabelNet's typologically diverse set of 499 world languages.
- 387, TITLE: On the Origin of Hallucinations in Conversational Models: Is it the Datasets or the Models?
<https://aclanthology.org/2022.naacl-main.387>
AUTHORS: Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane, Siva Reddy
HIGHLIGHT: Knowledge-grounded conversational models are known to suffer from producing factually invalid statements, a phenomenon commonly called hallucination. In this work, we investigate the underlying causes of this phenomenon: is hallucination due to the training data, or to the models?
- 388, TITLE: Is "My Favorite New Movie" My Favorite Movie? Probing the Understanding of Recursive Noun Phrases

- <https://aclanthology.org/2022.naacl-main.388>
AUTHORS: Qing Lyu, Zheng Hua, Daoxin Li, Li Zhang, Marianna Apidianaki, Chris Callison-Burch
HIGHLIGHT: We introduce the Recursive Noun Phrase Challenge (RNPC), a dataset of three textual inference tasks involving textual entailment and event plausibility comparison, precisely targeting the understanding of recursive NPs.
- 389, TITLE: Original or Translated? A Causal Analysis of the Impact of Translationese on Machine Translation Performance
<https://aclanthology.org/2022.naacl-main.389>
AUTHORS: Jingwei Ni, Zhijing Jin, Markus Freitag, Mrinmaya Sachan, Bernhard Schölkopf
HIGHLIGHT: In this work, we collect CausalMT, a dataset where the MT training data are also labeled with the human translation directions.
- 390, TITLE: Visual Commonsense in Pretrained Unimodal and Multimodal Models
<https://aclanthology.org/2022.naacl-main.390>
AUTHORS: Chenyu Zhang, Benjamin Van Durme, Zhuowan Li, Elias Stengel-Eskin
HIGHLIGHT: In this paper, we investigate to what degree unimodal (language-only) and multimodal (image and language) models capture a broad range of visually salient attributes.
- 391, TITLE: QuALITY: Question Answering with Long Input Texts, Yes!
<https://aclanthology.org/2022.naacl-main.391>
AUTHORS: Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, Samuel Bowman
HIGHLIGHT: To enable building and testing models on long-document comprehension, we introduce QuALITY, a multiple-choice QA dataset with context passages in English that have an average length of about 5,000 tokens, much longer than typical current models can process.
- 392, TITLE: ExSum: From Local Explanations to Model Understanding
<https://aclanthology.org/2022.naacl-main.392>
AUTHORS: Yilun Zhou, Marco Tulio Ribeiro, Julie Shah
HIGHLIGHT: In this paper, we introduce explanation summary (ExSum), a mathematical framework for quantifying model understanding, and propose metrics for its quality assessment.
- 393, TITLE: Maximum Bayes Smatch Ensemble Distillation for AMR Parsing
<https://aclanthology.org/2022.naacl-main.393>
AUTHORS: Young-Suk Lee, Ramn Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, Salim Roukos
HIGHLIGHT: However, for most recent high performant parsers, the effect of self-learning and silver data augmentation seems to be fading. In this paper we propose to overcome this diminishing returns of silver data by combining Smatch-based ensembling techniques with ensemble distillation.
- 394, TITLE: When Does Syntax Mediate Neural Language Model Performance? Evidence from Dropout Probes
<https://aclanthology.org/2022.naacl-main.394>
AUTHORS: Mycal Tucker, Tiwalayo Eisape, Peng Qian, Roger Levy, Julie Shah
HIGHLIGHT: We demonstrate that models do encode syntactic information redundantly and introduce a new probe design that guides probes to consider all syntactic information present in embeddings.
- 395, TITLE: Modeling Task Interactions in Document-Level Joint Entity and Relation Extraction
<https://aclanthology.org/2022.naacl-main.395>
AUTHORS: Liyan Xu, Jinho Choi
HIGHLIGHT: Especially, we address the two-way interaction between COREF and RE that has not been the focus by previous work, and propose to introduce explicit interaction namely Graph Compatibility (GC) that is specifically designed to leverage task characteristics, bridging decisions of two tasks for direct task interference.
- 396, TITLE: Few-Shot Semantic Parsing with Language Models Trained on Code
<https://aclanthology.org/2022.naacl-main.396>
AUTHORS: Richard Shin, Benjamin Van Durme
HIGHLIGHT: For semantic parsing tasks where we map natural language into code, such models may prove more adept at it. In this paper, we test this hypothesis and find that Codex performs better on such tasks than equivalent GPT-3 models.
- 397, TITLE: CORWA: A Citation-Oriented Related Work Annotation Dataset
<https://aclanthology.org/2022.naacl-main.397>
AUTHORS: Xiangci Li, Biswadip Mandal, Jessica Ouyang
HIGHLIGHT: As a first step toward a linguistically-motivated related work generation framework, we present a Citation Oriented Related Work Annotation (CORWA) dataset that labels different types of citation text fragments from different information sources.
- 398, TITLE: Overcoming Catastrophic Forgetting During Domain Adaptation of Seq2seq Language Generation
<https://aclanthology.org/2022.naacl-main.398>
AUTHORS: Dingcheng Li, Zheng Chen, Eunah Cho, Jie Hao, Xiaohu Liu, Fan Xing, Chenlei Guo, Yang Liu
HIGHLIGHT: In this work, we propose an innovative framework, RMR_DSE that leverages a recall optimization mechanism to selectively memorize important parameters of previous tasks via regularization, and uses a domain drift estimation algorithm to compensate the drift between different do-mains in the embedding space.

- 399, TITLE: Extreme Zero-Shot Learning for Extreme Text Classification
<https://aclanthology.org/2022.naacl-main.399>
AUTHORS: Yuanhao Xiong, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Inderjit Dhillon
HIGHLIGHT: In this paper, we consider a more practical scenario called Extreme Zero-Shot XMC (EZ-XMC), in which no supervision is needed and merely raw text of instances and labels are accessible.
- 400, TITLE: ConflIBERT: A Pre-trained Language Model for Political Conflict and Violence
<https://aclanthology.org/2022.naacl-main.400>
AUTHORS: Yibo Hu, MohammadSaleh Hosseini, Erick Skorupa Parolin, Javier Osorio, Latifur Khan, Patrick Brandt, Vito D'Orazio
HIGHLIGHT: To help advance research in political science, we introduce ConflIBERT, a domain-specific pre-trained language model for conflict and political violence.
- 401, TITLE: Automatic Multi-Label Prompting: Simple and Interpretable Few-Shot Classification
<https://aclanthology.org/2022.naacl-main.401>
AUTHORS: Han Wang, Canwen Xu, Julian McAuley
HIGHLIGHT: In this paper, we propose Automatic Multi-Label Prompting (AMuLaP), a simple yet effective method to automatically select label mappings for few-shot text classification with prompting.
- 402, TITLE: Few-shot Subgoal Planning with Language Models
<https://aclanthology.org/2022.naacl-main.402>
AUTHORS: Lajanugen Logeswaran, Yao Fu, Moontae Lee, Honglak Lee
HIGHLIGHT: Pre-trained language models have shown successful progress in many text understanding benchmarks. This work explores the capability of these models to predict actionable plans in real-world environments.
- 403, TITLE: IDPG: An Instance-Dependent Prompt Generation Method
<https://aclanthology.org/2022.naacl-main.403>
AUTHORS: Zhuofeng Wu, Sinong Wang, Jiatao Gu, Rui Hou, Yuxiao Dong, V.G.Vinod Vydiswaran, Hao Ma
HIGHLIGHT: In this paper, we propose a conditional prompt generation method to generate prompts for each input instance, referred to as the Instance-Dependent Prompt Generation (IDPG).
- 404, TITLE: Embedding Hallucination for Few-shot Language Fine-tuning
<https://aclanthology.org/2022.naacl-main.404>
AUTHORS: Yiren Jian, Chongyang Gao, Soroush Vosoughi
HIGHLIGHT: In this paper, we propose an Embedding Hallucination (EmbedHalluc) method, which generates auxiliary embedding-label pairs to expand the fine-tuning dataset.
- 405, TITLE: Cryptocurrency Bubble Detection: A New Stock Market Dataset, Financial Task & Hyperbolic Models
<https://aclanthology.org/2022.naacl-main.405>
AUTHORS: Ramit Sawhney, Shivam Agarwal, Vivek Mittal, Paolo Rosso, Vikram Nanda, Sudheer Chava
HIGHLIGHT: Taking the first step towards NLP for cryptocurrencies, we present and publicly release CryptoBubbles, a novel multi-span identification task for bubble detection, and a dataset of more than 400 cryptocurrencies from 9 exchanges over five years spanning over two million tweets.
- 406, TITLE: Nearest Neighbor Knowledge Distillation for Neural Machine Translation
<https://aclanthology.org/2022.naacl-main.406>
AUTHORS: Zhixian Yang, Renliang Sun, Xiaojun Wan
HIGHLIGHT: In this paper, we propose to move the time-consuming kNN search forward to the preprocessing phase, and then introduce k Nearest Neighbor Knowledge Distillation (kNN-KD) that trains the base NMT model to directly learn the knowledge of kNN.
- 407, TITLE: DEMix Layers: Disentangling Domains for Modular Language Modeling
<https://aclanthology.org/2022.naacl-main.407>
AUTHORS: Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah Smith, Luke Zettlemoyer
HIGHLIGHT: We introduce a new domain expert mixture (DEMix) layer that enables conditioning a language model (LM) on the domain of the input text.
- 408, TITLE: Contrastive Learning for Prompt-based Few-shot Language Learners
<https://aclanthology.org/2022.naacl-main.408>
AUTHORS: Yiren Jian, Chongyang Gao, Soroush Vosoughi
HIGHLIGHT: The impressive performance of GPT-3 using natural language prompts and in-context learning has inspired work on better fine-tuning of moderately-sized models under this paradigm. Following this line of work, we present a contrastive learning framework that clusters inputs from the same class for better generality of models trained with only limited examples.
- 409, TITLE: Cross-Lingual Event Detection via Optimized Adversarial Training
<https://aclanthology.org/2022.naacl-main.409>
AUTHORS: Luis Guzman-Nateras, Minh Van Nguyen, Thien Nguyen
HIGHLIGHT: In this work, we focus on Cross-Lingual Event Detection where a model is trained on data from a `\textit{source}` language but its performance is evaluated on data from a second, `\textit{target}`, language.

- 410, TITLE: Identifying Implicitly Abusive Remarks about Identity Groups using a Linguistically Informed Approach
<https://aclanthology.org/2022.naacl-main.410>
AUTHORS: Michael Wiegand, Elisabeth Eder, Josef Ruppenhofer
HIGHLIGHT: Following the recently-proposed strategy to solve implicit abuse by separately addressing its different subtypes, we present a new focused and less biased dataset that consists of the subtype of atomic negative sentences about identity groups.
- 411, TITLE: Label Definitions Improve Semantic Role Labeling
<https://aclanthology.org/2022.naacl-main.411>
AUTHORS: Li Zhang, Ishan Jindal, Yunyao Li
HIGHLIGHT: Learning symbolic labels usually requires ample training data, which is frequently unavailable due to the cost of annotation. We instead propose to retrieve and leverage the definitions of these labels from the annotation guidelines.
- 412, TITLE: Shedding New Light on the Language of the Dark Web
<https://aclanthology.org/2022.naacl-main.412>
AUTHORS: Youngjin Jin, Eugene Jang, Yongjae Lee, Seungwon Shin, Jin-Woo Chung
HIGHLIGHT: This paper introduces CoDA, a publicly available Dark Web dataset consisting of 10000 web documents tailored towards text-based Dark Web analysis.
- 413, TITLE: Conceptualizing Treatment Leakage in Text-based Causal Inference
<https://aclanthology.org/2022.naacl-main.413>
AUTHORS: Adel Daoud, Connor Jerzak, Richard Johansson
HIGHLIGHT: In this article, we define the treatment-leakage problem, and discuss the identification as well as the estimation challenges it raises.
- 414, TITLE: Consistency Training with Virtual Adversarial Discrete Perturbation
<https://aclanthology.org/2022.naacl-main.414>
AUTHORS: Jungsoo Park, Gyuwan Kim, Jaewoo Kang
HIGHLIGHT: Thus, the perturbed samples may not aid in regularization due to their ease of classification from the model. In this context, we propose an augmentation method of adding a discrete noise that would incur the highest divergence between predictions.
- 415, TITLE: CONFIT: Toward Faithful Dialogue Summarization with Linguistically-Informed Contrastive Fine-tuning
<https://aclanthology.org/2022.naacl-main.415>
AUTHORS: Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, Dragomir Radev
HIGHLIGHT: To tackle top factual errors from our annotation, we introduce additional contrastive loss with carefully designed hard negative samples and self-supervised dialogue-specific loss to capture the key information between speakers.
- 416, TITLE: CoMPM: Context Modeling with Speaker's Pre-trained Memory Tracking for Emotion Recognition in Conversation
<https://aclanthology.org/2022.naacl-main.416>
AUTHORS: Joosung Lee, Woojin Lee
HIGHLIGHT: We introduce CoMPM, which combines the speaker's pre-trained memory with the context model, and find that the pre-trained memory significantly improves the performance of the context model.
- 417, TITLE: Investigating Crowdsourcing Protocols for Evaluating the Factual Consistency of Summaries
<https://aclanthology.org/2022.naacl-main.417>
AUTHORS: Xiangru Tang, Alexander Fabbri, Haoran Li, Ziming Mao, Griffin Adams, Borui Wang, Asli Celikyilmaz, Yashar Mehdad, Dragomir Radev
HIGHLIGHT: To determine the factors that affect the reliability of the human evaluation, we crowdsource evaluations for factual consistency across state-of-the-art models on two news summarization datasets using the rating-based Likert Scale and ranking-based Best-Worst Scaling.
- 418, TITLE: DialSummEval: Revisiting Summarization Evaluation for Dialogues
<https://aclanthology.org/2022.naacl-main.418>
AUTHORS: Mingqi Gao, Xiaojun Wan
HIGHLIGHT: In our paper, we re-evaluate 18 categories of metrics in terms of four dimensions: coherence, consistency, fluency and relevance, as well as a unified human evaluation of various models for the first time.
- 419, TITLE: Hyperbolic Relevance Matching for Neural Keyphrase Extraction
<https://aclanthology.org/2022.naacl-main.419>
AUTHORS: Mingyang Song, Yi Feng, Liping Jing
HIGHLIGHT: Identifying important keyphrases is the central component of keyphrase extraction, and its main challenge is learning to represent information comprehensively and discriminate importance accurately. In this paper, to address the above issues, we design a new hyperbolic matching model (HyperMatch) to explore keyphrase extraction in hyperbolic space.
- 420, TITLE: Template-free Prompt Tuning for Few-shot NER
<https://aclanthology.org/2022.naacl-main.420>
AUTHORS: Ruotian Ma, Xin Zhou, Tao Gui, Yiding Tan, Linyang Li, Qi Zhang, Xuanjing Huang

- HIGHLIGHT: In this work, we propose a more elegant method to reformulate NER tasks as LM problems without any templates.
- 421, TITLE: Few-Shot Document-Level Relation Extraction
<https://aclanthology.org/2022.naacl-main.421>
AUTHORS: Nicholas Popovic, Michael F?rber
HIGHLIGHT: We present FREDo, a few-shot document-level relation extraction (FSDLRE) benchmark.
- 422, TITLE: LaMemo: Language Modeling with Look-Ahead Memory
<https://aclanthology.org/2022.naacl-main.422>
AUTHORS: Haozhe Ji, Rongsheng Zhang, Zhenyu Yang, Zhipeng Hu, Minlie Huang
HIGHLIGHT: As a result, this prohibits the memory to dynamically interact with the current context that provides up-to-date information for token prediction. To remedy this issue, we propose Look-Ahead Memory (LaMemo) that enhances the recurrence memory by incrementally attending to the right-side tokens and interpolating with the old memory states to maintain long-term information in the history.
- 423, TITLE: Exploiting Inductive Bias in Transformers for Unsupervised Disentanglement of Syntax and Semantics with VAEs
<https://aclanthology.org/2022.naacl-main.423>
AUTHORS: Ghazi Felhi, Joseph Roux, Djam? Seddah
HIGHLIGHT: We propose a generative model for text generation, which exhibits disentangled latent representations of syntax and semantics.
- 424, TITLE: Neighbors Are Not Strangers: Improving Non-Autoregressive Translation under Low-Frequency Lexical Constraints
<https://aclanthology.org/2022.naacl-main.424>
AUTHORS: Chun Zeng, Jiangjie Chen, Tianyi Zhuang, Rui Xu, Hao Yang, Qin Ying, Shimin Tao, Yanghua Xiao
HIGHLIGHT: In this paper, we focus on non-autoregressive translation (NAT) for this problem for its efficiency advantage.
- 425, TITLE: What do Toothbrushes do in the Kitchen? How Transformers Think our World is Structured
<https://aclanthology.org/2022.naacl-main.425>
AUTHORS: Alexander Henlein, Alexander Mehler
HIGHLIGHT: In this paper we utilize this research on biases to investigate to what extent transformer-based language models allow for extracting knowledge about object relations (X occurs in Y; X consists of Z; action A involves using X).
- 426, TITLE: Less is More: Learning to Refine Dialogue History for Personalized Dialogue Generation
<https://aclanthology.org/2022.naacl-main.426>
AUTHORS: Hanxun Zhong, Zhicheng Dou, Yutao Zhu, Hongjin Qian, Ji-Rong Wen
HIGHLIGHT: In this work, we propose to refine the user dialogue history on a large scale, based on which we can handle more dialogue history and obtain more abundant and accurate persona information.
- 427, TITLE: A Holistic Framework for Analyzing the COVID-19 Vaccine Debate
<https://aclanthology.org/2022.naacl-main.427>
AUTHORS: Maria Pacheco, Tunazzina Islam, Monal Mahajan, Andrey Shor, Ming Yin, Lyle Ungar, Dan Goldwasser
HIGHLIGHT: In this work we propose a holistic analysis framework connecting stance and reason analysis, and fine-grained entity level moral sentiment analysis.
- 428, TITLE: Learning to Win Lottery Tickets in BERT Transfer via Task-agnostic Mask Training
<https://aclanthology.org/2022.naacl-main.428>
AUTHORS: Yuanxin Liu, Fandong Meng, Zheng Lin, Peng Fu, Yanan Cao, Weiping Wang, Jie Zhou
HIGHLIGHT: These subnetworks are found using magnitude-based pruning. In this paper, we find that the BERT subnetworks have even more potential than these studies have shown.
- 429, TITLE: You Don't Know My Favorite Color: Preventing Dialogue Representations from Revealing Speakers' Private Personas
<https://aclanthology.org/2022.naacl-main.429>
AUTHORS: Haoran Li, Yangqiu Song, Lixin Fan
HIGHLIGHT: To this end, we propose effective defense objectives to protect persona leakage from hidden states.
- 430, TITLE: Explaining Dialogue Evaluation Metrics using Adversarial Behavioral Analysis
<https://aclanthology.org/2022.naacl-main.430>
AUTHORS: Baber Khalid, Sungjin Lee
HIGHLIGHT: In this work, we propose an adversarial test-suite which generates problematic variations of various dialogue aspects, e.g. logical entailment, using automatic heuristics.
- 431, TITLE: Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection
<https://aclanthology.org/2022.naacl-main.431>
AUTHORS: Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, Noah Smith
HIGHLIGHT: We seek to understand the *who*, *why*, and *what* behind biases in toxicity annotations.

- 432, TITLE: Non-Autoregressive Chinese ASR Error Correction with Phonological Training
<https://aclanthology.org/2022.naacl-main.432>
AUTHORS: Zheng Fang, Ruiqing Zhang, Zhongjun He, Hua Wu, Yanan Cao
HIGHLIGHT: As the errors introduced by ASR systems will impair the performance of downstream tasks, we introduce a post-processing error correction method, PhVEC, to correct errors in text space.
- 433, TITLE: Hate Speech and Counter Speech Detection: Conversational Context Does Matter
<https://aclanthology.org/2022.naacl-main.433>
AUTHORS: Xinchun Yu, Eduardo Blanco, Lingzi Hong
HIGHLIGHT: This paper investigates the role of context in the annotation and detection of online hate and counter speech, where context is defined as the preceding comment in a conversation thread.
- 434, TITLE: DACSA: A large-scale Dataset for Automatic summarization of Catalan and Spanish newspaper Articles
<https://aclanthology.org/2022.naacl-main.434>
AUTHORS: Encarnaci?n Segarra Soriano, Vicent Ahuir, Llu?s-F. Hurtado, Jos? Gonz?lez
HIGHLIGHT: In this work, we describe the construction of a corpus of Catalan and Spanish newspapers, the Dataset for Automatic summarization of Catalan and Spanish newspaper Articles (DACSA) corpus.
- 435, TITLE: Time Waits for No One! Analysis and Challenges of Temporal Misalignment
<https://aclanthology.org/2022.naacl-main.435>
AUTHORS: Kelvin Luu, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam, Noah Smith
HIGHLIGHT: In this work, we establish a suite of eight diverse tasks across different domains (social media, science papers, news, and reviews) and periods of time (spanning five years or more) to quantify the effects of temporal misalignment.
- 436, TITLE: MCSE: Multimodal Contrastive Learning of Sentence Embeddings
<https://aclanthology.org/2022.naacl-main.436>
AUTHORS: Miaoran Zhang, Marius Mosbach, David Adelani, Michael Hedderich, Dietrich Klakow
HIGHLIGHT: In this work, we propose a sentence embedding learning approach that exploits both visual and textual information via a multimodal contrastive objective.
- 437, TITLE: HiURE: Hierarchical Exemplar Contrastive Learning for Unsupervised Relation Extraction
<https://aclanthology.org/2022.naacl-main.437>
AUTHORS: Shuliang Liu, Xuming Hu, Chenwei Zhang, Shu?ang Li, Lijie Wen, Philip Yu
HIGHLIGHT: Existing works either utilize self-supervised schemes to refine relational feature signals by iteratively leveraging adaptive clustering and classification that provoke gradual drift problems, or adopt instance-wise contrastive learning which unreasonably pushes apart those sentence pairs that are semantically similar. To overcome these defects, we propose a novel contrastive learning framework named HiURE, which has the capability to derive hierarchical signals from relational feature space using cross hierarchy attention and effectively optimize relation representation of sentences under exemplar-wise contrastive learning.
- 438, TITLE: Diagnosing Vision-and-Language Navigation: What Really Matters
<https://aclanthology.org/2022.naacl-main.438>
AUTHORS: Wanrong Zhu, Yuankai Qi, Pradyumna Narayana, Kazoo Sone, Sugato Basu, Xin Wang, Qi Wu, Miguel Eckstein, William Yang Wang
HIGHLIGHT: In this work, we conduct a series of diagnostic experiments to unveil agents' focus during navigation.
- 439, TITLE: Aligning to Social Norms and Values in Interactive Narratives
<https://aclanthology.org/2022.naacl-main.439>
AUTHORS: Prithviraj Ammanabrolu, Liwei Jiang, Maarten Sap, Hannaneh Hajishirzi, Yejin Choi
HIGHLIGHT: We focus on creating agents that act in alignment with socially beneficial norms and values in interactive narratives or text-based games-environments wherein an agent perceives and interacts with a world through natural language.
- 440, TITLE: MOVER: Mask, Over-generate and Rank for Hyperbole Generation
<https://aclanthology.org/2022.naacl-main.440>
AUTHORS: Yunxiang Zhang, Xiaojun Wan
HIGHLIGHT: Despite being a common figure of speech, hyperbole is under-researched in Figurative Language Processing. In this paper, we tackle the challenging task of hyperbole generation to transfer a literal sentence into its hyperbolic paraphrase.
- 441, TITLE: Embarrassingly Simple Performance Prediction for Abductive Natural Language Inference
<https://aclanthology.org/2022.naacl-main.441>
AUTHORS: Emils Kadikis, Vaibhav Srivastav, Roman Klinger
HIGHLIGHT: This is a time-consuming and resource-intense endeavour. To solve this practical problem, we propose a simple method for predicting the performance without actually fine-tuning the model.
- 442, TITLE: Re-Examining System-Level Correlations of Automatic Summarization Evaluation Metrics
<https://aclanthology.org/2022.naacl-main.442>
AUTHORS: Daniel Deutsch, Rotem Dror, Dan Roth
HIGHLIGHT: We identify two ways in which the definition of the system-level correlation is inconsistent with how metrics are used to evaluate systems in practice and propose changes to rectify this disconnect.