Research Statement

Jing Jiang
School of Computing and Information Systems, Singapore Management University
Tel: (65) 6828-0785; Email: jingjiang@smu.edu.sg
18 December 2023

Background

My research revolves around Natural Language Processing, a key subfield of Artificial Intelligence that aims to understand human languages using computational methods. I have investigated broadly on the applied side of NLP, proposing new solutions based on principled machine learning models to problems in a range of areas including information extraction, topic modeling, sentiment analysis, social media analysis, question answering, and vision-language understanding.

A central concern that motivated my selection of research problems is that I see a prevalent and pressing need in real-world applications for advanced language technologies to quickly discover trends and patterns and to accurately extract knowledge from the huge amount of textual data surrounding us today. To address this pressing need, I have developed novel solutions to push the state of the art of language technologies.

Research Areas

My past and current research can be broadly organized into the following directions: question answering, out-of-distribution robustness, social media analysis, and information extraction.

Question Answering

Question answering (QA) is the task of answering natural language questions based on the content of documents, knowledge bases or other information sources. I have worked on both document-based QA and knowledge base QA (KBQA).

In [NAACL'16], we proposed an LSTM-based sequence matching model designed for natural language inference (NLI), which tries to determine whether a sentence entails, contradicts or is neutral to another sentence. After analyzing the nature of this sequence matching problem, we found that it was crucial to match word pairs from the two sequences, and we therefore designed a special Match-LSTM model which matches each word in the second sequence with an attention-weighted sum of tokens from the first sequence and then aggregates these matching results through an LSTM. This model performed significantly better than previous state of the art on the benchmark Stanford NLI dataset.

Subsequently, we saw the similarity between NLI and QA and applied Match-LSTM to the answer extraction problem in [ICLR'17a]. In addition to applying Match-LSTM, we also proposed to use a Pointer Network to predict the start and end positions of the answer span. We showed that our approach could substantially outperform previous methods on the benchmark SQuAD dataset. Meanwhile, realizing that

Match-LSTM can be generalized into a “compare-aggregate” framework, in [ICLR’17b] we systematically compared difference choices of matching functions on several different QA and NLI datasets to gain a better understanding of the compare-aggregate framework.

Following these three early pieces of work, we further investigated several other settings of the document-based QA task. In [ICLR’18] and [AAAI’18], we studied open-domain question answering and designed methods to aggregate evidence from different sources to strengthen question answering models. In [ACL’18], we proposed a co-matching model based on Match-LSTM to deal with three sequences for multiple choice QA.

More recently, my work on QA focused on knowledge base question answering. In [ASLP’19a], we extended the Match-LSTM framework for sequence matching to tackle the KBQA task. In [IJCAI’19a], we proposed a topic unit method to increase the coverage of candidate answers to improve KBQA performance. In [ICDM’19] and [ACL’20], we designed iterative query graph generation and pruning methods to handle the exponentially growing search space for multi-hop questions in KBQA. In [WSDM’21], we designed a teacher-student learning method to provide intermediate supervision signals for better reasoning in multi-hop KBQA. Our latest work in [ACL’21] explored conversational KBQA and proposed an entity transition graph to model the focal entities of a conversation to better encode a conversation history when answering a sequence of semantically related questions in a conversation. We conducted a comprehensive survey of recent work on KBQA in [IJCAI’21].

Currently my focus on question answering is on visual question answering (VQA). In this direction, I am interested in adopting pre-trained large-scale vision-language models for visual question answering when training data is limited. In [ACL’23 Findings], we proposed a modularized zero-shot network that explicitly decomposes questions into sub reasoning steps and is highly interpretable. We converted sub reasoning tasks to acceptable objectives of pre-trained models and assigned these sub reasoning tasks to proper pre-trained models without any adaptation. Our experiments on two VQA benchmarks under the zero-shot setting demonstrated the effectiveness of our method and better interpretability compared with several baselines.

While exploring VQA with pre-trained vision-language models, I realized that many state-of-the-art vision-language models still have their limitations. This inspired a related line of work that evaluates the capabilities of pre-trained vision-language models. In [AACL’22], we studied how to measure stereotypical bias in pre-trained vision-language models. We extended an existing text-only dataset, StereoSet, which covers a wide range of stereotypical bias, to a vision-language probing dataset called VLStereoSet. This dataset was designed for measuring stereotypical bias in vision-language models. We further defined several metrics to measure both a vision-language model’s overall stereotypical bias and its intra-modal and inter-modal bias. Experiments on six representative pre-trained vision-language models demonstrated that stereotypical biases clearly exist in most of these models. Subsequently, in [EMNLP’23 Findings], we introduced another novel probing dataset named ROME (reasoning beyond commonsense knowledge) to evaluate whether the state-of-the-art pre-trained vision-language models have the reasoning capability to correctly interpret counter-intuitive content. ROME contains images that defy commonsense

knowledge with regards to color, shape, material, size, and positional relation. Experiments on the state-of-the-art pre-trained vision-language models revealed that most of these models are still largely incapable of interpreting counter-intuitive scenarios.

Out-of-Distribution Robustness

A common problem faced by learning-based solutions to NLP problems is how to handle test data that comes from a distribution different from the training data. This includes scenarios when there is a domain shift (e.g., from news articles to online social media posts) or a language shift (e.g., from English to German) between the training and the test data. While this out-of-distribution problem has always been a challenge in NLP, it becomes more serious with recent deep learning models, which are good at "memorizing" training data, a phenomenon we examined in [ACL'22a].

My work to address the out-of-distribution problem started with domain adaptation, which I started investigating during my PhD years. My concern back then was to understand the nature of the problem and accordingly design general solutions. In [ACL'07], I conducted a theoretical analysis of the problem from a distributional point of view and identified two types of adaptation mechanisms, namely labeling adaptation and instance adaptation. I further proposed several heuristic instance weighting methods to achieve both kinds of adaptation. Subsequently in [CIKM'07], I developed another feature selection framework for domain adaptation. I introduced the notion of generalizable features and domain-specific features and proposed a two-stage approach consisting of generalization followed by adaptation.

With recent advances in deep neural network models, out-of-distribution (OOD) robustness has become an important topic in the NLP community. We recently developed an interventional training method to address the OOD problem [EMNLP'22]. Previous work using intervention to address OOD issues relies on known confounding biases to learn spurious correlations. We proposed a bottom-up automatic intervention method to automatically identify confounders and found the method to be effective on three benchmark datasets.

Another setting where test data differs from the training data is when the language changes, i.e., when we want to adapt a model trained on a source language for a different target language. A new approach to this cross-lingual transfer problem is to translate the training data from the source language to the target language using machine translation. However, machine-translated texts are often different from human-written texts in style, a phenomenon referred to as "translationese." In [ACL'22b], we address this translationese problem by explicitly modeling the differences between translated texts and original texts through a learned mapping function. Our method was evaluated on the multi-lingual QA dataset TyDiQA and demonstrated to be effective over strong baselines.

Next, I plan to further explore OOD robustness by leveraging large-scale pre-trained models, which presumably is robust on a wide range of domains because of the large coverage of their training data. The challenge we face is that these pre-trained models

are not directly trained for the downstream task, and fine-tuning or prompt-based learning is needed to adapt them for specific downstream tasks.

Social Media Analysis

I have worked on two major directions in social media analysis: modeling and understanding topics on social media, and mining sentiments and opinions on social media.

Social media provides platforms for people to easily express themselves and share information. Collectively these social media posts reflect the current interests and concerns of a population. Therefore, harvesting this topical information could be very useful to businesses and governments. Among the various social media platforms, Twitter is an open platform with a large user base. In 2011, we empirically compared the topical coverage of Twitter with traditional media [ECIR'11]. Several interesting findings were revealed by this work, including recognizing Twitter as a better source of entity-oriented topics than traditional media and the power of Twitter in spreading news but not in sharing opinions. In this work we also proposed a TwitterLDA model that was shown to be more suitable for tweets and was subsequently adopted in several other studies by other researchers. Given that topics on Twitter evolve quickly over time, I soon started investigating temporal modeling of Twitter topics to capture bursty topics that represent major short-term events. In [ACL'12], we designed a new topic model that automatically detects event-related tweets by distinguishing global time-dependent topic distributions from personal time-independent topic distributions. We further improved the model by modeling the association between bursty events and longstanding topics in [EMNLP'13] and by making the model non-parametric using recurrent Chinese restaurant process [SDM'14].

The other important aspect of social media content I looked at is its wide coverage of public sentiments and opinions. We proposed a novel MaxEnt-LDA hybrid model that combines supervised MaxEnt model with unsupervised LDA to separate words related to sentiment aspects and words related to sentiment polarities [EMNLP'10]. I have also explored cross-domain sentiment analysis. In [EMNLP'16], we used auxiliary tasks to help learn sentence embeddings that work well across domains to improve cross-domain sentiment classification. In [AAAI'17], we used auxiliary labels generated from syntactic rules that are less sensitive to domain changes to help improve an LSTM-based cross-domain opinion target extraction model. More recently, we investigated multi-modal sentiment analysis and leveraged image data to improve sentiment classification [IJCAI'19b, TASLP'20]. We also studied rumor detection and proposed a Transformer-based architecture for rumour detection and a coupled hierarchical Transformer model for stance-aware rumour verification [AAAI'20, EMNLP'20].

Information Extraction

To discover knowledge from unstructured text, an important first step is to extract information snippets, particularly entities and their relations. When I started investigating relation extraction, the start-of-the-art approach was supervised machine learning relying heavily on manual feature engineering. I proposed a unified graphical representation of the feature space for relation extraction, which enabled systematic exploration and comparison of different feature configurations to optimize

extraction performance [NAACL'07]. Another bottleneck in supervised approaches to information extraction is the need for sufficient labeled training data drawn from the same domain as the test data. In real applications, such labeled datasets are often not available. This challenge caught my attention back in 2005. I noticed that an entity recognizer trained on one domain often had degraded performance on a new domain because of domain differences. I subsequently designed a domain-aware feature ranking and selection strategy to address this problem [NAACL'06]. Later in [ACL'09], I studied a similar setting for relation extraction, where only a few labeled relation instances are available for a target relation type but plenty of training data can be borrowed from other relation types. I applied a domain adaptation framework I developed earlier and showed that the method successfully transfers common syntactic patterns across relation types, resulting in substantially improved performance.

Selected Publications and Outputs

[EMNLP'23 Findings] Kankan Zhou, Eason Lai, Wei Bin Au Yeong, Kyriakos Mouratidis and Jing Jiang, "ROME: Evaluating Pre-trained Vision-Language Models on Reasoning beyond Visual Common Sense," in Findings of the Association for Computational Linguistics: EMNLP, 2023.

[ACL'23 Findings] Rui Cao and Jing Jiang, "Modularized Zero-shot VQA with Pre-trained Models," in Findings of the Association for Computational Linguistics: ACL, 2023.

[AACL'22] Kankan Zhou, Eason Lai and Jing Jiang, "VLStereoSet: A Study of Stereotypical Bias in Pre-trained Vision-Language Models," in Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing, 2022.

[EMNLP'22] Sicheng Yu, Jing Jiang, Hao Zhang, Yulei Niu, Qianru Sun and Lidong Bing, "Interventional Training for Out-of-Distribution Natural Language Understanding," in Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2022.

[ACL'22a] Xiaosen Zheng and Jing Jiang, "An Emprical Study of Memorization in NLP," in Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, 2022.

[ACL'22b] Sicheng Yu, Qianru Sun, Hao Zhang and Jing Jiang, "Translate-Train Embracing Translationese Artifacts," in Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (short paper), 2022.

[IJCAI'21] Yunshi Lan, Gaole He, Jinhao Jiang, Jing Jiang, Wayne Xin Zhao and Ji-Rong Wen, "A survey on complex knowledge base question answering: Methods, challenges and solutions," in Proceedings of the 30th International Joint Conference on Artificial Intelligence (Survey Track), 2021.

[ACL'21] Yunshi Lan and Jing Jiang, "Modeling transitions of focal entities for conversational knowledge base question answering," in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics, 2021.

[WSDM'21] Gaole He, Yunshi Lan, Jing Jiang, Xin Zhao and Ji-Rong Wen, "Improving multi-hop knowledge base question answering by learning intermediate supervision signals," in Proceedings of the 14th ACM International WSDM Conference, 2021.

[ACL'20] Yunshi Lan and Jing Jiang, "Query graph generation for answering multi-hop complex questions from knowledge bases," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, (short paper), 2020.

[EMNLP'20] Jianfei Yu, Jing Jiang, Ling Min Serena Khoo, Hai Leong Chieu and Rui Xia, "Coupled hierarchical Transformer for stance-aware rumor verification in social media conversations," in Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, 2020.

[AAAI'20] Ling Min Serena Khoo, Hai Leong Chieu, Zhong Qian and Jing Jiang, "Interpretable rumor detection in microblogs by attending to user interactions," in Proceedings of the 34th AAAI Conference on Artificial Intelligence, 2020.

[TASLP'20] Jianfei Yu, Jing Jiang and Rui Xia, "Entity-sensitive attention and fusion network for entity-level multimodal sentiment classification," in IEEE/ACM Trans. Audio, Speech & Language Processing, 28, 429-439, 2020.

[IJCAI'19a] Yunshi Lan, Shuohang Wang and Jing Jiang, "Knowledge base question answering with topic units," in Proceedings of the 28th International Joint Conference on Artificial Intelligence, 2019.

[IJCAI'19b] Jianfei Yu and Jing Jiang, "Adapting BERT for target-oriented multimodal sentiment classification," in Proceedings of the 28th International Joint Conference on Artificial Intelligence, 2019.

[ICDM'19] Yunshi Lan, Shuohang Wang and Jing Jiang, "Multi-hop knowledge base question answering with an iterative sequence matching model," in Proceedings of the 19th IEEE International Conference on Data Mining, 2019.

[TASLP'19] Yunshi Lan, Shuohang Wang and Jing Jiang, "Knowledge base question answering with a matching-aggregation model and question-specific contextual relations," in IEEE/ACM Trans. Audio, Speech & Language Processing, 27 (10), 1629-1638, 2019.

[ACL'18] Shuohang Wang, Mo Yu, Shiyu Chang, Jing Jiang, "A co-matching model for multi-choice reading comprehension," in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, (short paper), 2018.

[AAAI'18] Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerald Tesauro, Bowen Zhou and Jing Jiang, "R3: Reinforced

ranker-reader for open domain question answering," in Proceedings of the 32nd AAAI Conference on Artificial Intelligence, 2018.

[ICLR'18] Shuohang Wang, Mo Yu, Jing Jiang,Wei Zhang, Xiaoxiao Guo, Shiyu Chang, Klinger. Tim, Gerald Tesauro and Murray Campbell, "Evidence aggregation for answer re-ranking in open domain question answering," in Proceedings of the 6th International Conference on Learning Representations, 2017.

[ICLR'17a] Shuohang Wang and Jing Jiang, "Machine comprehension using match-LSTM and answer pointer," in Proceedings of the 5th International Conference on Learning Representations, 2017.

[ICLR'17b] Shuohang Wang and Jing Jiang, "A compare-aggregate model for matching text sequences," in Proceedings of the 5th International Conference on Learning Representations, 2017.

[AAAI'17] Ying Ding, Jianfei Yu and Jing Jiang, "Recurrent neural networks with auxiliary labels for cross-domain opinion target extraction," in Proceedings of the 31st AAAI Conference on Artificial Intelligence, 2017.

[NAACL'16] Shuohang Wang and Jing Jiang, "Learning natural language inference with LSTM," in Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics, 2016.

[EMNLP'16] Jianfei Yu and Jing Jiang, "Learning sentence embeddings with auxiliary tasks for cross domain sentiment classification," in Proceedings of 2016 Conference on Empirical Methods in Natural Language Processing, 2016.

[SDM'14] Qiming Diao and Jing Jiang, "Recurrent Chinese restaurant process with a duration-based discount for event identification from Twitter," in Proceedings of the 2014 SIAM International Conference on Data Mining, 2014.

[EMNLP'13] Qiming Diao and Jing Jiang, "A unified model for topics, events and users on Twitter," in Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, 2013.

[ACL'12] Qiming Diao, Jing Jiang, Feida Zhu and Ee-Peng Lim, "Finding bursty topics from microblogs," in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics, 2012.

[ECIR'11] Wayne Xin Zhao, Jing Jiang, Jianshu Weng, Jing He, Ee-Peng Lim, Hongfei Yan and Xiaoming Li, "Comparing Twitter and traditional media using topic models," in Proceedings of the 33rd European Conference on Information Retrieval, 2011.

[EMNLP'10] Wayne Xin Zhao, Jing Jiang, Hongfei Yan and Xiaoming Li, "Jointly modeling aspects and opinions with a MaxEnt-LDA hybrid," in Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, 2010.

[ACL'09] Jing Jiang, "Multi-task transfer learning for weakly-supervised relation extraction," in Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics, 2009.

[CIKM'07] Jing Jiang and ChengXiang Zhai, "A two-stage approach to domain adaptation for statistical classifiers," in Proceedings of the ACM 16th Conference on Information and Knowledge Management, 2007.

[ACL'07] Jing Jiang and ChengXiang Zhai, "Instance weighting for domain adaptation in NLP," in Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, 2007.

[NAACL'07] Jing Jiang and ChengXiang Zhai, "A systematic exploration of the feature space for relation extraction," in Proceedings of the 2007 Conference of the North American Chapter of the Association for Computational Linguistics, 2007.

[NAACL'06] Jing Jiang and ChengXiang Zhai, "Exploiting domain structure for named entity recognition," in Proceedings of the 2006 Conference of the North American Chapter of the Association for Computational Linguistics, 2006.