

Jianfeng He

jianfenghe@vt.edu • 1-412-352-6807 • [Google Scholar](#) • [LinkedIn](#) • [Personal Website](#)

Research Areas: Trustworthy NLP (both text understanding and generation) via Uncertainty Estimation; Image Generation; Cross-modal Retrieval

EDUCATION

- Virginia Tech Falls Church, USA
Ph.D. Student in Computer Science *Jan. 2019 – Aug. 2024*
• Overall GPA: 3.95/4.0
- University of Pittsburgh Pittsburgh, USA
Ph.D. Student in Computer Science (Transferred to Virginia Tech) *Aug. 2018 – Dec. 2018*
• Admitted with first-year fellowship scholarship
- University of Chinese Academy of Science Beijing, China
Master in Computer Technology *Sep. 2014 – Jul. 2017*
• Admitted with exam exemption • Rank: Top 2
- Institute of Computing Technology, Chinese Academy of Sciences Beijing, China
Visiting Student in Group of Visual Information Processing and Learning *Jul. 2015 – Jul. 2017*
- Central China Normal University Wuhan, China
Bachelor in Digital Media Technology *Sep. 2010 – Jun. 2014*
• Rank: 2/43

WORK EXPERIENCE

- Amazon Development Center U.S., Inc. Manager: Hang Su
Applied Scientist Intern *May 2023 – Aug. 2023*
Project: Semi-Supervised Dialogue Abstractive Summarization via High-Quality Pseudolabel Selection.
Technology: Transformer, Text, Bayesian Neural Network, Beam Search Sampling.
- Amazon Development Center U.S., Inc. Manager: Kaisheng Yao
Applied Scientist Intern *May 2022 – Aug. 2022*
Project: Zero-Shot End-to-End Spoken Language Understanding via Cross-Modal Selective Self-Training.
Technology: Transformer, Text, Audio, Cluster, Self-training, Selective-learning.
- Tencent America, LLC Advisor: Linfeng Song
Natural Language & Speech Processing Research Intern *May 2021 – Aug. 2021*
Project: Pretrained Multimodal Dialogue System.
Technology: Transformer, Text, Image, Multimodal Fusion, Pre-trained Object Detection, OCR.

FIRST-AUTHOR PUBLICATIONS

- [1] (NAACL 2024) **J. He**, H. Su, J. Cai, H. Song, I. Igor. Semi-Supervised Dialogue Summarization via High-Quality Pseudolabel Selection. [C] Accept by NAACL 2024 main conference.
- [2] (NAACL 2024) **J. He**, L. Yu, S. Lei, C. Lu, F. Chen. Uncertainty Estimation on Sequential Labeling via Uncertainty Transmission. [C] Accept by NAACL 2024 Findings.
- [3] (EACL 2024) **J. He**, J. Salazar, K. Yao, H. Li, J. Cai. Zero-Shot End-to-End Spoken Language Understanding via Cross-Modal Selective Self-Training. [C] Accepted by EACL 2024 main conference.
- [4] (KDD 2023, research track) **J. He**, X. Zhang, S. Lei, F. Chen, A. Alhamadani, B. Xiao, C. Lu. CLUR: Uncertainty Estimation for Few-Shot Text Classification with Contrastive Learning. [C] In Proceedings of the 29th ACM SIGKDD international conference on knowledge discovery & data mining.
- [5] (SIGIR 2023, demo paper) **J. He**, S. Wu, A. Alhamadani, C. Chen, W. Lu, C. Lu, D. Solnick, Y. Li. Metro-Scope: An Advanced System for Real-Time Detection and Analysis of Metro-Related Threats and Events via Twitter. [C] In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval: 3130-3134.
- [6] (Neurocomputing 2022) **J. He**, X. Zhang, S. Lei, S. Wang, Q. Huang, C. Lu, B. Xiao. Semantic inpainting on segmentation map via multi-expansion loss. [J] Neurocomputing 501 (2022): 306-317.
- [7] (ICCV Workshop 2021) **J. He**, B. Xiao, X. Zhang, S. Wang, S. Lei, and C. Lu. Exploiting Characteristics in Semantic Inpainting on Segmentation Map: Semantic Metrics and Noise Reduction [C]. In Proceedings of the IEEE/CVF International Conference on Computer Vision 2021 (pp. 1876-1885).

- [8] (EMNLP 2020) **J. He**, X. Zhang, S. Lei, Z. Chen, F. Chen, A. Alhamadani, B. Xiao and C. Lu. [Towards More Accurate Uncertainty Estimation In Text Classification \[C\]](#). In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: 8362-8372.
- [9] (Neurocomputing 2018) **J. He**, B. Ma, S. Wang, Y. Liu and Q. Huang. [Multi-label Double Layers Learning for Cross-Modal Retrieval \[J\]](#). Neurocomputing, 2018, 275: 1893-1902.
- [10] (MATES Workshop 2017) **J. He**, S. Wang, Q. Qu, W. Zhang and Q. Huang. [Efficient Cross-Modal Retrieval Using Social Tag Information Towards Mobile Applications \[C\]](#). International Workshop on Mobility Analytics for Spatio-temporal and Social Data. Springer, Cham, 2017: 157-176.
- [11] (ACM Multimedia 2016, short paper) **J. He**, B. Ma, S. Wang, Y. Liu and Q. Huang. [Cross-modal Retrieval by Real Label Partial Least Squares \[C\]](#). Proceedings of the 2016 ACM on Multimedia Conference: 227-231.

FIRST-AUTHOR PREPRINTS

- [12] Y. Sun (co-first author), **J. He** (co-first author), L. Cui, S. Lei, C. Lu. [Med-MMHL: A Multi-Modal Dataset for Detecting Human- and LLM-Generated Misinformation in the Medical Domain](#). ArXiv preprint arXiv:2306.08871 (2023).

SELECTED CO-AUTHOR PUBLICATIONS

- [13] (ACL, 2023) S. Lei, X. Zhang, **J. He**, F. Chen, C. Lu. [TART: Improved Few-shot Text Classification Using Task-Adaptive Reference Transformation](#). [C] In 61st Annual Meeting of the Association for Computational Linguistics.
- [14] (ECCV, 2022) S. Lei, X. Zhang, **J. He**, F. Chen, C. Lu. [Cross-Domain Few-Shot Semantic Segmentation](#). [C] Cross-Domain Few-Shot Semantic Segmentation. In Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, pp. 73-90. Cham: Springer Nature Switzerland.
- [15] (NAACL Findings, 2022) S. Lei, X. Zhang, **J. He**, F. Chen, C. Lu. [Uncertainty-Aware Cross-Lingual Transfer with Pseudo Partial Labels](#). In Findings of the Association for Computational Linguistics: NAACL 2022 (pp. 1987-1997).
- [16] (ICME, 2021) S. Lei, X. Zhang, **J. He**, F. Chen, C. Lu. [Few-Shot Semantic Segmentation via Prototype Augmentation with Image-Level Annotations](#). In 2021 IEEE International Conference on Multimedia and Expo. 2021 Jul 5 (pp. 1-6). IEEE.
- [17] (IJCAI, 2017) L. Zhang, B. Ma, **J. He**, G. Li, Q. Huang and Q. Tian. [Adaptively Unified Semi-supervised Learning for Cross-Modal Retrieval \[C\]](#). Twenty-Sixth International Joint Conference on Artificial Intelligence, 2017: 3406-3412.

FIRST-AUTHOR ONGOING WORKS BEFORE GRADUATION

- [18] A Survey: [Uncertainty Estimation Models and Applications in Large Language Model](#).
- [19] A Benchmark for [Uncertainty Estimation in Large Language Model with Different Generation Metrics](#).

SELECTED AWARDS

Virginia Tech Computer Science Departmental **PhD research award**, Apr. 2024.
 SIGKDD **Student Travel Award** 2023, Jul. 2023.
 First-year **Fellowship Scholarship** Admission to U.Pitt, Aug. 2018 - May 2019
China National Scholarships (only awards for top 2% students), Sep. 2017
State-level First Prize of Chinese National Mathematical Modeling Contest, Sep. 2013

VOLUNTEER SERVICE

Local

- Job Fair Volunteer Department of Computer Science, Virginia Tech 2022
- Student Mentor for Graduate Student 2021

International

- ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL) 2021
- ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2023

PROFESSIONAL SERVICE

Paper Review

- Association for the Advancement of Artificial Intelligence (AAAI) 2021
- Conference on Computer Vision and Pattern Recognition (CVPR) 2021, 2022
- International Conference on Computer Vision (ICCV) 2021
- ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2022, 2023, 2024
- ACM Transactions on Knowledge Discovery from Data (TKDD) 2022
- IET Generation, Transmission & Distribution 2022
- Neural Information Processing Systems (NIPS-Datasets and Benchmarks) 2022, 2023
- The British Machine Vision Conference (BMVC) 2022, 2023
- ACL Rolling Review (ARR) 2022, 2023
- Empirical Methods in Natural Language Processing (EMNLP) 2023
- SIAM Conference on Data Mining (SDM) 2024

Research Talk

- Will give a talk “Uncertainty Estimation on Natural Language Processing” in June 2024 for a group in Amazon Search

TECHNICAL SKILLS

Languages and Technologies: C++, Matlab, Python(including **PyTorch**), Shell, Github, LaTeX.

Traditional Machine Learning Models: Subspace Learning, Metric Learning, Reinforcement Learning, etc.

Deep Learning Theory and Models: CNN, GCN, GAN, Transformer, BNN, ENN, Contrastive Learning, etc.

RESEARCH PROJECTS

Uncertainty Estimation In Dialogue Summarization

Advisor: **Hang Su**

Key Researcher

May 2023 – Aug. 2023

Project Goal: Develop uncertainty estimation models to assess the quality of generated text summaries.

My Responsibilities: Develop models to enhance the accuracy of uncertainty scores for generated texts.

My Achievements: Our proposed SiCF scores are effective on both uncertainty estimation and semi-supervised dialogue summarization. We plan to submit a paper for ARR OCT 2023.

Key Techniques: Bayesian Neural Network (BNN), Transformer, Beam Search Sampling.

Uncertainty Estimation In Natural Language Understanding

Advisor: **Chang-Tien Lu**

Key Researcher

Jan. 2019 – present

Project Goal: Develop uncertainty estimation models to analyze misclassification and out-of-domain samples in text classification, few-shot text classification, and named entity recognition.

My Responsibilities: Develop models to enhance the accuracy of uncertainty scores for both document-level and entity-level classified results.

My Achievements: Enhancing confidence in uncertainty scores for text classification by considering three distinct uncertainty types (aleatoric, epistemic, and parametric uncertainty) and mitigating overconfidence in top scores. Employing uncertainty relations for adaptive learning of uncertainty scores. Published two papers in EMNLP and KDD, with experiment-finished work planned for ARR OCT 2023.

Key Techniques: Data augmentation, Few-Shot, CNN, RNN, BNN, Transformer, Contrastive Learning.

Zero-Shot Spoken Language Understanding

Advisor: **Kaisheng Yao**

Key Researcher

May 2022 – Aug. 2022

Project Goal: Leveraging audio-text and text-semantics pairs to train a spoken language understanding model without requiring any audio-semantics pairs.

My Responsibilities: Develop a multi-modal to address the zero-shot spoken language understanding model and benchmark the task.

My Achievements: Significantly enhanced zero-shot spoken language understanding performance compared to baselines by introducing cross-modal selective self-training (CMSST) to tackle sample imbalance and label noise. One paper has been archived.

Key Techniques: Cluster, Self-training, Selective-learning.

Multimodal Dialogue System

Advisor: [Linfeng Song](#)
May 2021 – Aug. 2021

Key Researcher

Project Goal: Develop a multimodal dialogue system capable of generating system responses by considering the current user query and historical context, which can include both text and image inputs.

My Responsibilities: Develop and enhance the performance of a multi-modal dialogue system.

My Achievements: On 1% data, enhanced the performance of the multi-modal dialogue system compared to baselines by leveraging more fine-grained image information, including objects and their attributes, as well as incorporating OCR information from images.

Key Techniques: Multimodal Fusion, Transformer, Pre-trained Object Detection, Pre-trained OCR.

Semantic Inpainting On Segmentation Maps (SISM)

Advisor: [Chang-Tien Lu](#)
Aug. 2019 – May 2021

Key Researcher

Project Goal: Develop and enhance SISM models that focus on inpainting a masked area within a segmentation map based on the semantics defined by a target label, prioritizing coherence with the intended context.

My Responsibilities: Proposed more effective models to address the limitations of current global and local GANs in solving SISM, leveraging the unique characteristics present in segmentation maps.

My Achievements: Enhanced result consistency through the introduction of a novel multi-expansion loss, resulting in a significant reduction of noise pixels in the inpainted segmentation maps. Developed a novel metric to evaluate the semantic quality of the inpainted segmentation maps. Published two papers in ICCVW and Neurocomputing, respectively.

Key Techniques: GAN, Conditional GAN, Dilation Operation, and Image Processing.

Cross-Modal Retrieval

Advisors: [Bingpeng Ma](#) & [Qingming Huang](#)
Aug. 2015 – Feb. 2017

Key Researcher

Project Goal: Develop a model capable of retrieving results in one modality based on content similarity using queries from another modality.

My Responsibilities: Enhance the model performance in both single-label and multi-label settings.

My Achievements: Significant advancements have been achieved in cross-modal retrieval, including the development of a two-layer subspace learning model that iteratively learns a shared space for multiple modalities and latent spaces for multi-label labels. Additionally, a novel label representation has been introduced for single-label scenarios by integrating it into KPLS. Three papers have been published in ACM Multimedia, Neurocomputing, and a workshop.

Key Techniques: Dimension Reduction (PLS, KPLS), Matrix Derivations, Non-Linear Projection.