Haoyuan Li

https://leehaoyuan.github.io/ | haoyuanl@cs.unc.edu | https://scholar.google.com/citations?hl=en&user=BdfQ44MAAAAJ

EDUCATION

University of North Carolina at Chapel Hill, Chapel Hill, US Sep.2022-Present
Major: Computer Science, Doctor of Philosophy, Department of Computer Science
- **Advisor: Snigdha Chaturvedi**

Harvard University, Cambridge, US Sep.2020-Jan.2022
Major: Health Data Science, Master of Science, School of Public Health
- **Cumulative GPA:** 3.79/4.0

Renmin University of China, Beijing, CN Sep.2016–July.2020
Major: Statistics, Bachelor of Science, School of Statistics
- **Cumulative GPA:** 3.74/4.0

INTERNSHIP

Amazon May.2025-Aug.2025
- **Position:** Applied Scientist Intern **Mentor:** Zhengyuan Shen

Microsoft Research Asia Jan.2021-July.2021
- **Position:** Research Intern **Mentor:** He Di

Yale University Jun.2020-Sep.2020
- **Position:** Research Assistant **Advisor:** Dragomir R. Radev

Stanford University Jun.2019-Sep.2019
- **Position:** Research Assistant **Advisor:** Ram Rajagopal

PUBLICATIONS

Haoyuan Li, Zhengyuan Shen, Sullam Jeoung, Yueyan Chen, Jiayu Li, Qi Zhu, Shuai Wang, Vassilis N. Ioannidis, Huzefa Rangwala, "BoundRL: Efficient Structured Text Segmentation through Reinforced Boundary Generation", submitted to ICLR 2026

Haoyuan Li, Snigdha Chaturvedi, "Comparative Personalization for Multi-document Summarization", Submitted to ACL ARR July

Haoyuan Li, Rui Zhang, Snigdha Chaturvedi, 'Improving Fairness of Large Language Models in Multi-document Summarization', ACL 2025

Haoyuan Li, Yusen Zhang, Rui Zhang, Snigdha Chaturvedi, 'Coverage-based Fairness in Multi-document Summarization', NAACL 2025 SAC award for Summarization

Haoyuan Li, Snigdha Chaturvedi, 'Rationale-based Opinion Summarization', NAACL 2024

Haoyuan Li, Somnath Basu Roy Chowdhury, Snigdha Chaturvedi, 'Aspect-aware Unsupervised Extractive Opinion Summarization', ACL 2023 Findings

Bi, Qiwei, **Haoyuan Li**, and Hanfang Yang. 2021. "Boosting Few-Shot Abstractive Summarization with Auxiliary Tasks.", CIKM 2021

Fabbri, Alexander R., Simeng Han, **Haoyuan Li**, Haoran Li, Marjan Ghazvininejad, Shafiq Joty, Dragomir Radev, and Yashar Mehdad. "Improving Zero and Few-Shot Abstractive Summarization with Intermediate Fine-tuning and Data Augmentation.", NAACL 2021

Wang, Z.*, **Haoyuan Li***, Rajagopal, R., (2019) Urban2Vec: Incorporating Street View Imagery and POIs for Multi-Modal Urban Neighborhood Embedding, AAAI 2020, (*contribute equally)

REASEARCH & PROJECTS

BoundRL: Efficient Structured Text Segmentation through Reinforced Boundary Generation

Advisor: Shen Zhengyuan, Amazon

Summary: We propose BoundRL, a novel and efficient approach that jointly performs token-level text segmentation and label prediction for long structured texts. Instead of generating complete contents for each segment, it generates only a sequence of starting tokens, To adapt the model for the output format, BoundRL performs reinforcement learning with verifiable rewards (RLVR) with a reward that jointly optimizes document reconstruction fidelity and semantic alignment. To mitigate entropy collapse, it further constructs intermediate candidates by systematically perturbing a fraction of generated sequences of segments to create stepping stones toward higher-quality solutions.

Comparative Personalization for Multi-document Summarization

Advisor: Snigdha Chaturvedi, University of North Carolina at Chapel Hill

Summary: We propose ComPSum, a personalized MDS framework. It first generates a structured analysis of a user by comparing their preferences with other users' preferences. The generated structured analysis is then used to guide the generation of personalized summaries. To evaluate the performance of ComPSum, we propose AuthorMap, a fine-grained reference-free evaluation framework for personalized MDS. It evaluates the personalization of a system based on the authorship attribution between two personalized summaries generated for different users.

Improving Fairness of Large Language Models in Multi-document Summarization

Advisor: Snigdha Chaturvedi, University of North Carolina at Chapel Hill

Summary: To improve the fairness of summaries generated by LLMs, we propose FairPO, a preference tuning method based on DPO. To generate preference pair for tuning, the summaries are generated based on perturbed input document sets. To further improve the corpus-level fairness, we dynamically adjust the weights of each preference pairs.

Coverage-based Fairness in Multi-Document
Advisor: Snigdha Chaturvedi, University of North Carolina at Chapel Hill
Summary: To measure the fairness of summaries generated by LLMs, we propose a new summary-level fairness measure, Equal Coverage based on coverage of documents with different social attribute values. Equal Coverage addresses the redundancy issue faced by previous measures like Proportional Representation. We also propose a new corpus-level measure, Coverage Parity, which measures the corpus-level fairness in addition to Equal Coverage.

Rationale-based Opinion Summarization
Advisor: Snigdha Chaturvedi, University of North Carolina at Chapel Hill
Summary: To improve the usefulness and interpretability of opinion summaries, we propose a new paradigm, Rationale-based Opinion Summarization, where each represented opinion is paired with a rationale that provides more details. To extract rationales for a representative opinion, we extract sentences based on four criteria: relatedness, specificity, popularity, and diversity.

Aspect-aware Unsupervised Extractive Opinion Summarization
Advisor: Snigdha Chaturvedi, University of North Carolina at Chapel Hill
Summary: To ensure coverage of diverse aspects, we introduce an unsupervised extractive opinion summarization method that automatically identifies the aspects described in the review sentences and then extracts sentences based on their aspects. It identifies the underlying aspects of the review sentences using the roots of noun phrases and adjectives appearing in them. It then extracts review sentences of different aspects to form the summary.

Boosting few-shot abstractive summarization with auxiliary tasks
Advisor: Hanfang Yang, Renmin University of China
Summary: To improve the sample efficiency, we develop 3 auxiliary tasks for each training sample in addition to the main task: sentence extraction, object prediction, and triple entailment. We then trained the abstractive summarization model with these tasks in a multi-task learning setting. In this project, I build the BART model for abstractive summarization, use adapters for multi-task learning, run experiments on Xsum, and write the paper.

Improving Zero and Few-Shot Abstractive Summarization with Intermediate Fine-tuning and Data Augmentation
Advisor: Dragomir R. Radev, Yale University
Summary: To enable domain adaptation of summarization system with a few labeled samples, we generate pseudo-training pairs from the wiki dataset based on the characteristics of the target dataset such as abstractiveness. We also introduce back-translation and consistency loss to further improve the performance. In the project, I introduce consistency loss between original and augmented documents given pseudo-target summaries and write the paper.

Augmented Abstractive Summarization with Document-Level Semantic Graph
Advisor: Hanfang Yang, Asso Prof. School of Statistics, Renmin University of China
Summary: To improve the performance of news summarization, we generate a graph of entities in each document using distant supervision to better model the relationship between entities. Then, we apply GNN to generate the entity embedding and integrate them into the summarization model using an attention mechanism. In the project, I construct the graph and write the paper.

Urban2Vec: Incorporating Street View Imagery and POIs for Multi-Modal Urban Neighborhood Embedding
Advisor: Ram Rajagopal, Stanford University
Summary: To facilitate real estate recommendation and demographic surveys, we develop unsupervised methods to generate the embedding of neighborhoods based on the street view image and points of interest. To evaluate the effectiveness of our method, we use the embeddings to predict the demographic features of neighborhoods and run clustering on the embeddings. In the project, I design a way to combine image and text and design methods to evaluate its performance.