

# Ruochen (Esther) Zhao

+65 84322044 | ruochen002@e.ntu.edu.sg

## EDUCATION

---

### Nanyang Technological University

*Ph.D. in Artificial Intelligence; GPA: 4.5 out of 5.0*

Mentored by Prof. Shafiq Joty Rayhan, member of NTU-NLP team

Scholarship: One of the 4 recipients for the AI.SG Ph.D. scholarship in 2021 Spring

Research focus: Faithfulness and Reasoning in Large Language Models, Explaining complex NLP Systems

Singapore

Jan. 2021 – Present

### Harvard University

*M.S. in Data Science; GPA: 3.8 out of 4.0*

Relevant Coursework: Machine Learning, Time Series, Data Science I and II, Business Analytics, Statistical Inference

Course Projects: Spotify Recommender System, Inventory Optimization using Gurobi, Predicting Food Reviews with RNN in Machine Learning, Image Modification using CNN and Bayesian Autoencoders

Cambridge, MA

Aug. 2019 – Dec. 2020

### New York University

*B.A. in Mathematics; Double Minor in Computer Science and Economics; GPA: 3.8 out of 4.0*

Honors: Dean's Award, Mathematics Award, National Phi Beta Kappa Honors Society, Dean's List Honors for four years

Fellowships: Women In Science Fellowship, Grace Hopper Fellowship

New York, NY

Aug. 2015 – May. 2019

## PUBLICATIONS

---

### Chain of Knowledge: A Framework for Grounding Large Language Models with Structured Knowledge Bases

Xingxuan Li\*, Ruochen Zhao\*, Yew Ken Chia\*, Bosheng Ding, Lidong Bing, Shafiq Joty, Soujanya Poria

- Proposes a post-hoc interpretability method for extracting predictive high-level features (concepts) from the pretrained model's hidden layer activations. Extensive experiments on real and synthetic tasks demonstrate that our method achieves superior results on predictive impact, usability, and faithfulness compared to the baselines.

In Submission

### Explaining Language Models' Predictions with High-Impact Concepts

Ruochen Zhao, Shafiq Joty, Yongjie Wang, Prathyusha Jwalapuram, Tan Wang

- Proposes a post-hoc interpretability method for extracting predictive high-level features (concepts) from the pretrained model's hidden layer activations. Extensive experiments on real and synthetic tasks demonstrate that our method achieves superior results on predictive impact, usability, and faithfulness compared to the baselines.

In Submission

### PromptSum: Planning with Mixed Prompts for Parameter-Efficient Controllable Abstractive Summarization

Mathieu Ravaut\*, Hailin Chen\*, Ruochen Zhao\* (co-first author), Chengwei Qin\*, Shafiq Joty, Nancy F Chen

- Introduces PromptSum, a method combining Prefix Tuning with a multi-task objective and discrete entity prompts for abstractive summarization. Our model achieves state-of-the-art results on several popular few-shot benchmarks as well as a strong level of controllability through entities, all while only tuning several orders of magnitude less parameters.

In Submission

### Retrieving Multimodal Information for Augmented Generation: A Survey

Ruochen Zhao, Hailin Chen, Weishi Wang, Fangkai Jiao, Xuan Long Do, Chengwei Qin, Bosheng Ding, Xiaobao Guo, Minzhi Li, Xingxuan Li, Shafiq Joty

- Reviews methods that retrieve multimodal knowledge to assist and augment generative models, which offer a promising solution to important concerns such as factuality, reasoning, interpretability, and robustness.

In Submission

### Verify-and-Edit: A Knowledge-Enhanced Chain-of-Thought Framework

Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, Lidong Bing

- Combines Chain of Thought (CoT) with knowledge retrieval to enhance factuality of LLMs. Building on top of GPT3, the framework leads to accuracy improvements on multiple open-domain question-answering tasks.

ACL' 23

### Randomized Smoothing with Masked Inference for Adversarially Robust NLP Systems

Han Cheol Moon, Shafiq Joty, Ruochen Zhao, Megh Thakkar, Xu Chi

- Introduces RSMI, a novel two-stage framework that combines randomized smoothing (RS) with masked inference (MI) to improve the adversarial robustness of NLP systems.
- Contributed in writing the rigorous mathematical proof (Pg 14 – Pg 18)

ACL' 23

## WORK EXPERIENCE

---

### Rex

*Summer Data Scientist*

**Redwood City, CA**

*Jun. 2020 – Aug. 2020*

- Led a Natural Language Generation project in automatically generating house sale descriptions by learning human-written descriptions, used Encoder-Decoder Network with Attention, and trained on a large amount of data using GPU
- Used feed-forward neural network with self-defined loss to build an interpretable neural network for price prediction
- Analyzed the effect of holding house showings virtually on housing transactions using statistical hypothesis testing coded in Python, used SQL queries to pull and clean data from the company database, and presented the final report to stakeholders
- Company presentation on autoencoders, utilized git workflow management, and developed production-ready packages

### Compstak. Inc.

*Summer Data Analyst*

**New York, NY**

*May. 2018 – Aug. 2018*

- Processed and analyzed more than 10,000 computational data in the crowdsourced commercial real estate industry through formula-based excel work and internal data pipeline, including data on rent, transaction SQFT, building amenities, etc
- Utilized internal data pipeline (SpecialK) to specialize in Utah Real Estate market trends and improved its grouping
- Conducted data cleaning and data standardization projects to optimize data transparency

### Sugih Energy International Ltd

*Summer Finance Analyst*

**Singapore**

*Jun. 2017 – Aug. 2017*

- Conducted detailed industry research by gathering data from institutional reports on market dynamics, competitive landscape, and dynamic coal pricing on the international energy market; presented research briefings to seniors in the group
- Gathered data and developed a quantitative analysis of the pricing trends from China, Australia, and Indonesia by building Excel models to help the firm make informed decisions regarding future deal flows of coal international trade

## RESEARCH PROJECTS

---

### World Health Organization

*Research Assistant*

**Singapore**

*Mar. 2022 – Jul. 2022*

- Devised NLP pipeline to help WHO in automatically analyzing a large amount of COVID-related articles
- Contributed in designing the topic modeling component, which helped in drafting multiple COVID reports
- Collaborated with Lee Kong Chian School of Medicine, WHO, and NTU-NLP lab

### Harvard University

*Research Assistant*

**Cambridge, MA**

*Aug. 2020 - Now*

- Conducted research with Dr. David Perkes in the Econ-CS lab while focusing on Natural Language Processing
- Used BERT and tf-idf models to build predictive price models from email text data

### New York University

*Independent Researcher*

**New York, NY**

*Sep. 2018 – Dec. 2018*

- Mentored by a senior professor in the NYU Math for Finance graduate program to lead research on Statistical Arbitrage
- Analyzed ~1,000,000 stock data, formed pairs using 3 cointegration tests, empirically tested trading strategies by coding up a simulation trading in Python, and critically analyzed its strengths, weaknesses, and real-life applications

## SKILLS

---

- **Programming Skills:** Proficient in Python, Java, SQL, Excel, Powerpoint; Experienced in Pytorch and Tensorflow machine learning development; Intermediate in Matlab, R, C; Relevant projects in NLP and CV
- **Interests:** Texas Hold'em Poker, Travel blog, Kayaking, and Canoeing