

JUSTIN CHIH-YAO CHEN

✉ cychen@cs.unc.edu [Website](#) [LinkedIn](#) [Google Scholar](#) [Citations: 1,039](#) [h-index: 13](#)

EDUCATION



The University of North Carolina at Chapel Hill

Ph.D. student in Computer Science

Chapel Hill, NC

Aug. 2023 – Present



National Taiwan University

Master of Data Science. GPA: 4.10 / 4.30

Taipei, Taiwan

Sep. 2020 – Jun. 2022



National Cheng Kung University

Bachelor of Statistics, Double Major in Computer Science. GPA: 3.93 / 4.30

Tainan, Taiwan

Sep. 2016 – Jun. 2020

EXPERIENCE



MURGe-Lab NLP Group

Research Assistant (Advisor: [Mohit Bansal](#))

Improving LLM exploration, efficiency, and multi-agent collaboration for reasoning

Chapel Hill, NC

Aug. 2023 – Present



Microsoft Research

Research Intern (Mentors: [Harsh Jhamtani](#), [Gaurav Verma](#), [Benjamin Van Durme](#))

Learning self-evolving agent skills from user log data

Redmond, WA

May 2026 – Aug. 2026



Salesforce AI Research

Research Intern (Mentors: [Xiangyu Peng](#), [Prafulla Kumar Choubey](#), [Jason Wu](#))

Pushing the boundaries of LLM reasoning via post-training with self-generated hints

Palo Alto, CA

May 2025 – Aug. 2025



Google Cloud AI Research

Student Researcher (Mentors: [Chen-Yu Lee](#), [Zifeng Wang](#))

Enhancing the reasoning of LLMs by strengthening their reverse-thinking capabilities

Sunnyvale, CA

Jun. 2024 – Sep. 2024



Institute of Information Science, Academia Sinica

Research Assistant (Advisor: [Lun-Wei Ku](#))

Working on attractive headline generation, representation learning, and vision question generation

Taipei, Taiwan

Jul. 2020 – Dec. 2022



Institute of Data Science, National Cheng Kung University

Research Assistant (Advisor: [Cheng-Te Li](#))

Working on zero-shot relation extraction and anomaly detection

Tainan, Taiwan

Jun. 2018 – Jun. 2020

PUBLICATIONS (*CO-FIRST AUTHOR)

- [1] **Justin Chih-Yao Chen***, Sukwon Yun*, Elias Stengel-Eskin*, Tianlong Chen, Mohit Bansal. “Skill-Based Mixture-of-Experts: Adaptive Routing for Heterogeneous Reasoning via Inferred Skills.” **ICML 2026**. [\[Paper\]](#) [\[Code\]](#)
- [2] **Justin Chih-Yao Chen**, Becky Xiangyu Peng, Prafulla Kumar Choubey, Kung-Hsiang Huang, Jiaxin Zhang, Mohit Bansal, Chien-Sheng Wu. “Nudging the Boundaries of LLM Reasoning.” **ICLR 2026**. [\[Paper\]](#) [\[Code\]](#)
- [3] **Justin Chih-Yao Chen**, Zifeng Wang, Hamid Palangi, Rujun Han, Sayna Ebrahimi, Long Le, Vincent Perot, Swaroop Mishra, Mohit Bansal, Chen-Yu Lee, Tomas Pfister. “Reverse Thinking Makes LLMs Stronger Reasoners.” **NAACL 2025**. [\[Paper\]](#) [\[Code\]](#)
- [4] **Justin Chih-Yao Chen**, Archiki Prasad, Swarnadeep Saha, Elias Stengel-Eskin, Mohit Bansal. “MAgICoRe: Multi-Agent, Iterative, Coarse-to-Fine Refinement for Reasoning.” **EMNLP 2025**. [\[Paper\]](#) [\[Code\]](#)
- [5] **Justin Chih-Yao Chen**, Swarnadeep Saha, Mohit Bansal. “ReConcile: Round-Table Conference Improves Reasoning via Consensus among Diverse LLMs.” **ACL 2024**. [\[Paper\]](#) [\[Code\]](#)
- [6] **Justin Chih-Yao Chen***, Swarnadeep Saha*, Elias Stengel-Eskin, Mohit Bansal. “MAGDI: Structured Distillation of Multi-Agent Interaction Graphs Improves Reasoning in Smaller Language Models.” **ICML 2024**. [\[Paper\]](#) [\[Code\]](#)

- [7] **Justin Chih-Yao Chen**, Tun-Min Hung, Lun-Wei Ku. “Label-Aware Hyperbolic Embeddings for Fine-grained Emotion Classification.” **ACL 2023**. [\[Paper\]](#) [\[Code\]](#)
- [8] **Justin Chih-Yao Chen**, Dennis Wu, Lun-Wei Ku. “HonestBait: Forward References for Attractive but Faithful Headline Generation.” **ACL Findings 2023**. [\[Paper\]](#)
- [9] **Justin Chih-Yao Chen** and Cheng-Te Li. “ZS-BERT: Towards Zero-Shot Relation Extraction with Attribute Representation Learning.” **NAACL 2021**. [\[Paper\]](#) [\[Code\]](#)
- [10] **Justin Chih-Yao Chen**, Archiki Prasad, Zaid Khan, Joykirat Singh, Runchu Tian, Elias Stengel-Eskin, Mohit Bansal. “Cog-DRIFT: Exploration on Adaptively Reformulated Instances Enables Learning from Hard Reasoning Problems.” arXiv preprint. [\[Paper\]](#) [\[Code\]](#)
- [11] Hyunji Lee*, **Justin Chih-Yao Chen***, Joykirat Singh, Zaid Khan, Elias Stengel-Eskin, Mohit Bansal. “MINTEval: Evaluating Memory under Multi-Target Interference in Long-Horizon Agent Systems.” arXiv preprint. [\[Paper\]](#) [\[Code\]](#)
- [12] Jaewoo Lee, Archiki Prasad, **Justin Chih-Yao Chen**, Zaid Khan, Elias Stengel-Eskin, Mohit Bansal. “PRInTS: Reward Modeling for Long-Horizon Information Seeking.” **ACL 2026**.
- [13] Joykirat Singh, **Justin Chih-Yao Chen**, Archiki Prasad, Elias Stengel-Eskin, Akshay Nambi, Mohit Bansal. “Think Right: Learning to Mitigate Under-Over Thinking via Adaptive, Attentive Compression.” **COLM 2026**.
- [14] David Wan, **Justin Chih-Yao Chen**, Elias Stengel-Eskin, Mohit Bansal. “MAMM-Refine: A Recipe for Improving Faithfulness in Generation with Multi-Agent Collaboration.” **NAACL 2025**.
- [15] Swarnadeep Saha, Archiki Prasad, **Justin Chih-Yao Chen**, Peter Hase, Elias Stengel-Eskin, Mohit Bansal. “System-1.x: Learning to Balance Fast and Slow Planning with Language Models.” **ICLR 2025**.

HONORS & AWARDS

College Student Research Creativity Award (Out of 2,689 Research)	2020
<i>Ministry of Science and Technology, Taiwan</i>	
Honorable Mention Award, Big Data and Business Analysis Competition	2018
<i>Department of Finance, National Sun Yat-sen University</i>	
Cathay United Bank Undergraduate Scholarship	2017
<i>Cathay Life Insurance Company</i>	
Academic Excellence Award (Top 10%)	2016
<i>Department of Statistics, National Cheng Kung University</i>	

REFERENCES

Mohit Bansal Director of the MURGe-Lab, UNC Chapel Hill mbansal@cs.unc.edu	<i>PhD advisor</i>	Jason Wu Senior Director, Salesforce AI Research wu.jason@salesforce.com	<i>Internship mentor</i>
Chen-Yu Lee Principal Scientist and Director, Google Cloud AI Research chenyulee@google.com	<i>Internship mentor</i>	Hamid Palangi Staff Research Scientist, Google; Associate Professor, UW hamidpalangi@google.com	<i>Internship mentor</i>
Zifeng Wang Research Scientist, Google Cloud AI Research zifengwangking@gmail.com	<i>Internship mentor</i>	Harsh Jhamtani Machine Learning Researcher, Apple harshjhamtani@gmail.com	<i>Internship mentor</i>
Gaurav Verma Applied Scientist, IDEAS Research, Microsoft gauravverma@microsoft.com	<i>Internship mentor</i>	Benjamin Van Durme Professor, JHU; Research Manager, Microsoft vandurme@jhu.edu	<i>Internship mentor</i>
Lun-Wei Ku Director of the NLP/SA Lab, Institute of Information Science, Academia Sinica lwku@iis.sinica.edu.tw	<i>Master's advisor</i>	Cheng-Te Li Director of the Networked AI Lab, National Cheng Kung University chengte@ncku.edu.tw	<i>Undergraduate advisor</i>