JIACHENG XU
Research Scientist
Salesforce AI Research
LinkedIn Email: jcxu17@gmail.com
https://jiacheng-xu.github.io/

(a) Intro

I am a research scientist specializing in NLP. I have been leading the data and evaluation efforts for code LLMs in product, and have conducted research on topics including LLM evaluation, text decoding algorithms, SWE agents, etc. Currently, I work on improving code LLMs with verified rewards, and manage code annotators as well as partnerships with external data vendors.

(b) Education

The University of Texas at Austin	Computer Science	Doctorate Program , 2017 – 2022
Fudan University	Computer Science	B.E. *summa cum laude*, 2013 – 2017

(c) Professional Experience

Salesforce AI Research Palo Alto, CA
Senior Research Scientist Feb 2024 – Present
Research Scientist Jun 2022 – Jan 2024
Managers: Yingbo Zhou, Caiming Xiong

- **Core Contributor to AgentForce for Developers:** Key member of the team developing Salesforce's customer-facing Code LLM product (e.g., code autocompletion, chat, bug fixing), powered by in-house large language models.
  - **Team Management:** Recruited and managed a team of 4 contract code annotators. Served as co-lead for major data partnerships with Scale AI (2023) and Turing (2024 – present).
  - **Data Lead:** Oversaw data collection pipeline, including web data crawling, data augmentation, code analysis with custom parsers, regulatory compliance (e.g. PII removal), creation of annotation guidelines and data specifications, and quality assurance of annotated data.
  - **Evaluation Lead:** Designed and implemented automatic execution-based benchmarks and human evaluation pipelines; led model evaluation and selection for production release; extended benchmarks to support new features.
- **Research Contributions:** Published 7 papers and filled 5 patents since I joined Salesforce in top venues like COLM, ICLR, and ACL, including 2 as first author: neural text generation (Xu et al., 2023) and AI for software development (automated test case generation; Xu et al., 2025).
  - **Research interest**: long chain-of-actions for code LLMs.
  - **Mentorship:** Supervised PhD research interns: Renze Lou (2024), Jipeng Zhang (2025), Zeyu Leo Liu (2025, upcoming).

Google New York, NY
Researcher intern at Google Assistant, working on disfluency detection
Mentor: Shyam Upadhyay, Manaal Faruqui, Aditya Gupta & Diyi Yang, May – August 2020

Microsoft Bellevue, WA
Researcher intern at Microsoft Dynamics 365 AI Research, working on text summarization
Mentor: Zhe Gan, Yu Cheng & Jingjing Liu, June – August 2019

(d) Academic Experience

The University of Texas at Austin Austin, TX
Graduate research assistant at TAUR Lab
Advisor: Greg Durrett, August 2017 – May 2022
Thesis Committee: Greg Durrett, Ray Mooney, Junyi Li, Kyunghyun Cho (NYU)

Fudan University Shanghai, China
Undergraduate research assistant at FudanNLP
Thesis: *Long Text Encoding via Recurrent Neural Networks*
Advisor: Xipeng Qiu & Xuanjing Huang, June 2015 – February 2017

(e) Publications

Recent Publications

1. Jiacheng Xu, Bo Pang, Jin Qu, Hiroaki Hayashi, Caiming Xiong, and Yingbo Zhou, Clover: A test case generation benchmark with coverage, long-context, and verification (2025), arXiv:2502.08806 [cs.SE] .
2. Bo Pang, Hanze Dong, Jiacheng Xu, Silvio Savarese, Yingbo Zhou, and Caiming Xiong, Bolt: Bootstrap long chain-of-thought in language models without distillation, arXiv preprint arXiv:2502.03860 (2025).
3. Jiacheng Xu, Caiming Xiong, Silvio Savarese, and Yingbo Zhou, Best-k search algorithm for neural text generation, in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (2023) pp. 12385–12401.
4. Kexun Zhang, Weiran Yao, Zuxin Liu, Yihao Feng, Zhiwei Liu, Rithesh R N, Tian Lan, Lei Li, Renze Lou, Jiacheng Xu, Bo Pang, Yingbo Zhou, Shelby Heinecke, Silvio Savarese, Huan Wang, and Caiming Xiong, Diversity empowers intelligence: Integrating expertise of software engineering agents, in *The Thirteenth International Conference on Learning Representations* (2025).
5. Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg Durrett, A long way to go: Investigating length correlations in rlhf, in *First Conference on Language Modeling*.
6. Prasann Singhal, Jiacheng Xu, Xi Ye, and Greg Durrett, Eel: Efficiently encoding lattices for reranking, in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (2023) pp. 9299–9316.
7. Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, *et al.*, Beyond the imitation game: Quantifying and extrapolating the capabilities of language models, Transactions on Machine Learning Research (2023).

Text Summarization

5. Jiacheng Xu, Siddhartha Jonnalagadda, and Greg Durrett, Massive-scale decoding for text generation using lattices, in *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2022) pp. 4659–4676.
6. Jiacheng Xu and Greg Durrett, Dissecting generation modes for abstractive summarization models via ablation and attribution, in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural*

Language Processing (Volume 1: Long Papers) (Association for Computational Linguistics, Online, 2021) pp. 6925–6940.

7. Jiacheng Xu, Shrey Desai, and Greg Durrett, Understanding neural abstractive summarization models via uncertainty, in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Online, 2020) pp. 6275–6281.
8. Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu, Discourse-aware neural extractive text summarization, in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Association for Computational Linguistics, Online, 2020) pp. 5021–5031.
9. Jiacheng Xu and Greg Durrett, Neural extractive text summarization with syntactic compression, in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (Association for Computational Linguistics, Hong Kong, China, 2019) pp. 3292–3303.
10. Liyan Tang, Tanya Goyal, Alex Fabbri, Philippe Laban, Jiacheng Xu, Semih Yavuz, Wojciech Kryściński, Justin Rousseau, and Greg Durrett, Understanding factual errors in summarization: Errors, summarizers, datasets, error detectors, in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (2023) pp. 11626–11644.
11. Tanya Goyal, Jiacheng Xu, Junyi Jessy Li, and Greg Durrett, Training dynamics for text summarization models, in *Findings of the Association for Computational Linguistics: ACL 2022* (2022) pp. 2061–2073.
12. Ojas Ahuja, Jiacheng Xu, Akshay Gupta, Kevin Horecka, and Greg Durrett, Aspect-oriented summarization through query-focused extraction (2021), arXiv:2110.08296 [cs.CL] .
13. Shrey Desai, Jiacheng Xu, and Greg Durrett, Compressive summarization with plausibility and salience modeling, in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Online, 2020) pp. 6259–6274.

Other Publications

14. Jiacheng Xu and Greg Durrett, Spherical latent spaces for stable variational autoencoders, in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, Brussels, Belgium, 2018) pp. 4503–4513.
15. Jiacheng Xu, Xipeng Qiu, Kan Chen, and Xuanjing Huang, Knowledge graph representation with jointly structural and textual encoding, in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17* (2017) pp. 1318–1324.
16. Jiacheng Xu, Danlu Chen, Xipeng Qiu, and Xuanjing Huang, Cached long short-term memory neural networks for document-level sentiment classification, in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, Austin, Texas, 2016) pp. 1660–1669.
17. Lifu Tu, Semih Yavuz, Jin Qu, Jiacheng Xu, Rui Meng, Caiming Xiong, and Yingbo Zhou, Unlocking anticipatory text generation: A constrained approach for faithful decoding with large language models, arXiv preprint arXiv:2312.06149 (2023).
18. Aditya Gupta, Jiacheng Xu, Shyam Upadhyay, Diyi Yang, and Manaal Faruqui, Disfl-QA: A benchmark dataset for understanding disfluencies in question answering, in *Findings of the*

Association for Computational Linguistics: ACL-IJCNLP 2021 (Association for Computational Linguistics, Online, 2021) pp. 3309–3319.

19. Greg Durrett, Jifan Chen, Shrey Desai, Tanya Goyal, Lucas Kabela, Yasumasa Onoe, and Jiacheng Xu, Contemporary NLP modeling in six comprehensive programming assignments, in *Proceedings of the Fifth Workshop on Teaching NLP* (Association for Computational Linguistics, Online, 2021) pp. 99–103.
20. Jinyue Su, Jiacheng Xu, Xipeng Qiu, and Xuanjing Huang, Incorporating discriminator in sentence generation: a gibbs sampling method, in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence* (2018).

(f) Awards

1. Outstanding reviewer for ACL-2021, 2021
2. 2021 Global Top 100 Chinese Rising Stars in AI (25 NLP recipients worldwide, 0.2%), Baidu Scholar & Baidu Talent Intelligence Center, 2021
3. Travel grant for Amazon's Graduate Research Symposium, 2019
4. Outstanding undergraduate thesis award (1 per department), Fudan University, 2017
5. EMC Research Scholarship, EMC Lab (China), 2016

(g) Professional Service

1. Conference Program Committee member (reviewer): AAAI (2020, 2021, 2022), NeurIPS (2020, 2021), ICML(2021), ICLR (2021, 2022), ACL (2020, 2021, 2022), EMNLP (2021), COLING (2020), CoNLL (2020), NAACL-HLT (2019, 2021), CCL (2020)
2. Conference Program Committee member (chair): Publication chair of Diversity & Inclusion (D&I) Committee, NAACL-HLT 2022
3. Conference Workshop Program Committee member (reviewer): DeepLo (EMNLP 2019 Workshop), SRW (NAACL-HLT 2019 Workshop), NUSE (ACL 2020 Workshop), SRW (ACL-IJCNLP 2021 Workshop), NeuSum (EMNLP 2021 Workshop)
4. Journal Reviewer: Journal of Artificial Intelligence Research (JAIR) (2020, 2021), Computational Linguistics (CL) (2020), ACM Transactions on Asian and Low-Resource Language Information Processing (2020)

(h) Departmental Service

1. Member of PhD Junior Admissions Committee, UT Austin, 2020 – 2021
2. Mentor of Graduate Application Assistance Program, UT Austin, 2020 – 2021
3. Member of Masters Admissions Committee, UT Austin, 2018 – 2019, 2019 – 2020

(i) Teaching Experience

Teaching Assistant CS378: Natural Language Processing (undergraduate), UT Austin, January 2019 – May 2019

(j) Mentoring Experience

1. Ojas Ahuja (2020 – 2021, undergraduate at UT Austin, now at Waymo)
2. Shrey Desai (2019 – 2020, undergraduate at UT Austin, now at Facebook)
3. Akshay Gupta (2019 – 2020, undergraduate at UT Austin, now at Plaid)