Ben (Tong) Niu

mrnt0810@gmail.com linkedin.com/in/tong-niu-b1910658 Google Scholar

Summary

AI researcher with 9+ years in large language models, retrieval, and natural language generation, published at ACL, EMNLP, NAACL, ICLR, AAAI, and TACL, with embedding and reranking models ranked #1 on MTEB and BRIGHT. Takes research into production: a unified on-device speech model at Archer Aviation, and an LLM-generated regression test suite for the Salesforce Core app that catches release regressions before customers do, saving millions of dollars in manual testing per release.

Experience

Staff AI Research Engineer, *Archer Aviation* 2025 – 2026

- Spearheaded the 0-to-1 design, training, and rollout of an end-to-end on-device speech model that replaced four legacy systems with higher generation quality, earning a discretionary executive RSU award for accelerated production impact
- Invented a timestamp-to-token alignment algorithm that removed the need for a forced aligner
- Optimized the model for on-device iOS inference and completed its full production deployment within latency targets
- Managed and mentored two direct reports, setting technical direction and guiding daily execution

Senior Research Scientist, *Meta Superintelligence Labs* 2025

- Built an OCR-augmented vision-language model pipeline that grounds homework answers in the recognized problem text
- Built the evaluation harness and error taxonomy separating OCR failures from reasoning failures to target fixes
- Improved homework-answer accuracy over the base vision-language model across question types

Senior Research Scientist, *Salesforce AI Research* 2019 – 2025

- **Test generation:** Generated a large regression test suite for the Salesforce Core app with LLMs, catching regressions in new releases before customers do and saving millions of dollars in manual testing per release
- **Reranking:** Built JudgeRank, an agentic LLM reranker for reasoning-intensive retrieval that held the #1 spot on the BRIGHT benchmark for six months
- **Embeddings:** Co-developed and served the SFR-Embedding model, ranked #1 on MTEB
- **Pretraining data:** Built a Selenium-based crawler for large-scale academic corpora used in LLM pretraining, improving MMLU performance
- **LLM safety:** Developed DetoxiGen, a parameter-efficient contrastive-decoding method that steers LLM output away from toxic content
- **Applied NLP:** Built HTML-aware machine translation for localization services, RAG-based question answering for banking regulatory compliance, and a recommendation system for email subject-line generation
- **Research:** Published at ACL, EMNLP, NAACL, ICLR, and ACM TOCHI on paraphrasing, cross-lingual retrieval, question generation, and dialogue evaluation
- **Internship:** Started as a Research Intern, building an unsupervised BERT-based sentence-compression model

Research Assistant, *University of North Carolina at Chapel Hill* (with Prof. Mohit Bansal) 2017 – 2019

- Adapted AutoAugment to NLP, learning data-augmentation policies for dialogue tasks (EMNLP 2019)
- Designed adversarial strategies exposing over-sensitive, over-stable dialogue models (CoNLL 2018)
- Built RL-based polite dialogue generation that requires no parallel data (TACL 2018)

Education

Duke University, M.S. in *Computer Science* (PhD program, PhD Fellowship) 2014 – 2016

University of Virginia, B.S. in *Mathematics* and *Computer Science*, minor in *Physics* (in three years) 2011 – 2014

Selected Publications

1. **T. Niu**, S. Joty, Y. Liu, C. Xiong, Y. Zhou, and S. Yavuz. *JudgeRank: Leveraging Large Language Models for Reasoning-Intensive Reranking*. **arXiv:2411.00142, 2024**.
2. **T. Niu**, C. Xiong, Y. Zhou, and S. Yavuz. *Parameter-Efficient Detoxification with Contrastive Decoding*. HuCLLM @ **ACL 2024**.
3. W. Zhao, Y. Liu, **T. Niu**, Y. Wan, P. Yu, S. Joty, Y. Zhou, and S. Yavuz. *DivKnowQA: Assessing the Reasoning Ability of LLMs via Open-Domain Question Answering over Knowledge Base and Text*. Findings of **NAACL 2024**.
4. X. Chen, C.-S. Wu, L. Murakhovs'ka, P. Laban, **T. Niu**, W. Liu, and C. Xiong. *Marvista: Exploring the Design of a Human-AI Collaborative News Reading Tool*. **ACM TOCHI 2023**.
5. L. Murakhovs'ka, C.-S. Wu, P. Laban, **T. Niu**, W. Liu, and C. Xiong. *MixQG: Neural Question Generation with Mixed Answer Types*. Findings of **NAACL 2022**.
6. **T. Niu**, K. Hashimoto, Y. Zhou, and C. Xiong. *OneAligner: Zero-Shot Cross-Lingual Transfer with One Rich-Resource Language Pair for Low-Resource Sentence Retrieval*. Findings of **ACL 2022**.
7. G. Aguilar, B. McCann, **T. Niu**, N. Rajani, N. S. Keskar, and T. Solorio. *Char2Subword: Extending the Subword Embedding Space Using Robust Character Compositionality*. Findings of **EMNLP 2021**.
8. S. Li, S. Yavuz, K. Hashimoto, J. Li, **T. Niu**, N. Rajani, X. Yan, Y. Zhou, and C. Xiong. *CoCo: Controllable Counterfactuals for Evaluating Dialogue State Trackers*. **ICLR 2021**.
9. **T. Niu**, S. Yavuz, Y. Zhou, N. S. Keskar, H. Wang, and C. Xiong. *Unsupervised Paraphrasing with Pretrained Language Models*. **EMNLP 2021**.
10. **T. Niu** and M. Bansal. *AvgOut: A Simple Output-Probability Measure to Eliminate Dull Responses*. **AAAI 2020**.
11. **T. Niu** and M. Bansal. *Automatically Learning Data Augmentation Policies for Dialogue Tasks*. **EMNLP 2019**.
12. S. Karlekar, **T. Niu**, and M. Bansal. *Detecting Linguistic Characteristics of Alzheimer's Dementia by Interpreting Neural Models*. **NAACL 2018**.
13. **T. Niu** and M. Bansal. *Adversarial Over-Sensitivity and Over-Stability Strategies for Dialogue Models*. **CoNLL 2018**.
14. **T. Niu** and M. Bansal. *Polite Dialogue Generation Without Parallel Data*. **TACL 2018**.

Skills

- **Research areas:** Large language models, text embeddings and reranking, retrieval-augmented generation, LLM agents, vision-language models, reinforcement learning, machine translation, speech modeling and alignment
- **Engineering:** Python, PyTorch, on-device (iOS) inference optimization, model serving, web crawling (Selenium)
- **Languages:** English (fluent), Chinese (native)

Awards

- **FY25 Team Astro Award**, Atlas, Salesforce AI Research 2024
- **Facebook ParlAI Research Award**, Facebook AI Research 2017
- **GPNANO Fellowship**, Nanoscience, Duke University 2015
- **PhD Fellowship**, Computer Science, Duke University 2014

Books & Ventures

- **Study Grammar Like a Pro**, a top-selling English grammar book written in Chinese 2013
- **Founder**, *Jake's American English*: founded and ran it before college, growing to 450 students and 2 branches 2010
- **Left and Right: The Conjoined Brothers**, world's first graded Chinese reader series (300-word level) 2008