

Donghong Cai

☎ 314-706-2000 | @ cai.d@wustl.edu | 🔗 LinkedIn | 🌐 <https://ilikevegetable.github.io>

I am a third-year Ph.D. student in Computer Science at Washington University in St. Louis, advised by Dr. Yixin Chen. My research focuses on diffusion language models, structured data generation, and LLM reasoning, with a background in self-supervised learning. I am interested in LLM, Agentic AI, reinforcement learning, and self-improving AI systems, and am seeking **Spring or Summer 2027 Research Scientist Intern and Applied Scientist Intern** opportunities.

EDUCATION

Washington University in St. Louis <i>Ph.D. in Computer Science; GPA: 3.96/4.00</i>	St. Louis, USA <i>Aug 2024 – present</i>
Zhejiang University <i>M.S. in Software Engineering; GPA: 3.84/4.00</i>	Hangzhou, China <i>Sep 2020 – Mar 2023</i>
Huazhong University of Science and Technology <i>B.S. in Software Engineering; GPA: 3.67/4.00</i>	Wuhan, China <i>Sep 2016 – Jun 2020</i>

WORK EXPERIENCE

Alibaba Group <i>Machine Learning Engineer, International E-commerce Seller Team</i>	Shenzhen, China <i>Apr 2023 – Jul 2024</i>
- Developed an LLM-based seller assistant with a hybrid retrieval framework, combining semantic matching over QA pairs for frequent queries with retrieval-augmented generation (RAG) over databases and documents. - Developed recommendation algorithms for merchant promotional tools, increasing merchant GMV by 8.28% and units ordered by 6.55% on average. - Built time-series models for monthly merchant GMV and order-volume forecasting, improving prediction accuracy by 21.7% and 36.8%, respectively, over manual heuristics. - Received a top-tier annual performance rating of 3.75 (top 10%) for FY2024 (Apr 2023–Mar 2024).	

PUBLICATIONS & PREPRINTS

* Equal contribution.

- MBrain: A Multi-channel Self-Supervised Learning Framework for Brain Signals**
Donghong Cai*, Junru Chen*, Yang Yang, Teng Liu, Yafeng Li.
KDD 2023
- TabDLM: Free-Form Tabular Data Generation via Joint Numerical–Language Diffusion**
Donghong Cai, Jiarui Feng, Yanbo Wang, Da Zheng, Yixin Chen, Muhan Zhang.
NeurIPS 2026
- Dependency-Structured Generation with Masked Diffusion Language Models**
Hanqing Zeng*, Diyi Hu*, Donghong Cai*, Jiarui Feng*, Yizhe Zhang, Yinglong Xia, Yangtian Zhang, Spilios Evmorfos, Jiayi Liu, Lizhu Zhang, Hongwei Wang, Hongyu Xiong, Yang Song, Xiangjun Fan, Yixin Chen, Fei Sha.
Preprint, 2026
- GRIP: In-Parameter Graph Reasoning through Fine-Tuning Large Language Models**
Jiarui Feng, Donghong Cai, Yixin Chen, Muhan Zhang.
KDD 2026
- Nonsense Helps: Prompt Space Perturbation Broadens Reasoning Exploration**
Langlin Huang, Chengsong Huang, Jinyuan Li, Donghong Cai, Yuyi Yang, Jiaxin Huang.
EMNLP 2026

6. Questioning the Questions: Sustaining Self-Evolution in Reasoning Models

Jinyuan Li, Chengsong Huang, Langlin Huang, **Donghong Cai**, Shiping Gao, Yuyi Yang, Jiaxin Huang.
Preprint, 2026

7. Process Rewards with Learned Reliability

Jinyuan Li, Langlin Huang, Chengsong Huang, Shaoyang Xu, **Donghong Cai**, Yuyi Yang, Wenxuan Zhang, Jiaxin Huang.
Preprint, 2026

SELECTED RESEARCH EXPERIENCE

TabDLM: Joint Numerical–Language Diffusion

NeurIPS 2026

- Developed a unified generator for numerical, categorical, and free-form text fields by integrating continuous numerical diffusion with masked language diffusion; designed numerical tokenization and LoRA adaptation of LLaDA-8B to capture dependencies across fields.
- Implemented joint training, sampling, and benchmark evaluation; reduced pairwise dependency error by **85.6% on Amazon** and **37.1% on Arxiv** relative to the strongest baseline on each dataset, while achieving the best downstream predictive utility.

MAPLE: Dependency-Structured Generation

Preprint, 2026

- Developed MAPLE to train masked diffusion LMs on teacher-extracted token-dependency DAGs; derived an ELBO objective that jointly learns reveal positions and token values, with data-distribution recovery guarantees under model realizability.
- Implemented continued pretraining of LLaDA 2.0-16B MoE and DreamReasoner-8B; achieved **69.26% average math accuracy** on DreamReasoner (**+2.63 percentage points** over the strongest baseline), while retaining standard parallel decoding without inference-time graphs.

GRIP: In-Parameter Graph Reasoning

KDD 2026

- Developed a parameter-efficient framework that stores graph knowledge in LoRA adapters for reasoning without graph context at inference; designed graph memorization, summarization, and question-answering tasks for fine-tuning Qwen2.5-7B and Llama3.1-8B.
- Built evaluation and ablation pipelines for knowledge-graph and compositional reasoning tasks; achieved **7.85× per-query inference speedup** on FB15K237 with Llama3.1-8B versus the graph-context baseline, after one-time adapter training.

TEACHING EXPERIENCE

Washington University in St. Louis

St. Louis, USA

Teaching Assistant

Fall 2026

- CSE 5403 – Algorithms for Nonlinear Optimization