CSE 8803 DLT: Deep Learning for Text Data (Fall 2026)

Logistics

- **Instructor**: Chao Zhang (chaozhang@gatech.edu)
- **Lecture time**: TBD
- **Location**: TBD
- **Office Hours**: TBD

Course Description

This course studies large language models (LLMs) and their applications in text data analysis, reasoning, and autonomous agents. The course covers LLM foundations (pretraining, alignment, inference), LLM reasoning (prompting strategies, verifiers, test-time compute scaling), and LLM agents (agent architectures, agent training, and evaluation). Each lecture session combines an instructor overview with student paper presentations and discussion. Students will also compete on a leaderboard challenge and build an LLM-based system as their course project.

Prerequisites: (1) basic knowledge in machine learning (e.g., CX 4240, CSE 6740, CS 4641); (2) comfortable reading research papers and giving presentations; (3) solid programming skills—the course project requires building working systems.

Course Topics

Part I: LLM Foundations

- Text representation: Word2Vec, Transformer, BERT, GPT
- Scaling laws and emergent abilities
- Instruction tuning, prompting, in-context learning, chain-of-thought
- LLM alignment: RLHF, DPO, GRPO
- LLM inference: quantization, KV caching, serving

Part II: LLM Reasoning

- Reasoning strategies: tree-of-thought, self-consistency, self-refinement
- Verifiers: process reward models vs outcome reward models
- Test-time compute scaling: the o1/o3 paradigm
- RL for reasoning: GRPO, STaR, reward shaping

Part III: LLM Agents

- Agent architecture and harness: ReAct, Reflexion, planning, memory, tool use
- Agent training: RL for agents, trajectory optimization, self-improvement
- Evaluation and benchmarks: SWE-bench, WebArena, OSWorld
- Frontiers: multi-agent systems, cross-domain applications

Course Policies

- **Academic Integrity**: Students are expected to abide by the Georgia Tech Honor Code.
- **Use of AI tools**: Students may use AI assistants for the leaderboard challenge and project. Paper reviews and presentations must reflect your own understanding.
- **Accommodations**: Students with learning needs requiring special accommodation should contact the Office of Disability Services as soon as possible.
- **Student-Faculty Expectations**: See catalog.gatech.edu/rules/22.

Resources

- Speech and Language Processing, by Jurafsky and Martin
- Deep Learning, by Goodfellow, Bengio, and Courville
- Dive into Deep Learning, by Zhang, Lipton, Li, and Smola
- Additional research papers provided throughout the course.

This syllabus is preliminary. The detailed schedule, paper list, and grading policy will be finalized at the start of the semester.