Haosheng Wang

Providence, RI, 02903 | 6469230226 | hswang947@gmail.com | https://wonder947.com

EDUCATION

Brown University, Providence, RI — *Sep 2024 – May 2026*
- *Master of Science in Computer Science* — *GPA: 4.0/4.0*

New York University, New York, NY — *Jan 2021 – May 2024*
- *Bachelor of Arts in Computer Science & Math* — *GPA: 3.905/4.0*
- *Honors: Dean's List for 3 Academic Years, Magna Cum Laude*

WORKING EXPERIENCE

Linksome, Shenzhen, China — *May 2025 – Aug 2025*

AI Algorithm Intern
- Designed and implemented an ontology-driven knowledge graph RAG system for large-scale domain-specific corpora.
- Developed an autonomous pipeline that jointly constructs a KG-grounded ontology and an ontology-aligned knowledge graph, enabling structured, database-style querying.
- Proposed two complementary information retrieval mechanisms: (1) an iterative agentic search strategy, and (2) a query-decomposed graph pattern matching approach, targeting complex multi-entity, multi-hop queries.
- Evaluated the system against contemporary graph-based RAG frameworks (e.g., GraphRAG, LightRAG) in terms of cost, latency, accuracy, recall, interpretability, and scalability.

RESEARCH EXPERIENCE

Knowledge Graph Construction — *May 2024 – Feb 2025*

NYU Independent Research, Advisor: Dr. Hasan Aljabbouli
- Analyzed current research landscape of knowledge-graph construction (KGC) from unstructured texts, and categorized them into ontology-driven, ontology-evolving, and non-ontology KGC paradigms.
- Identified critical evaluation metrics for each KGC paradigm such as Precision, Recall, F1, hallucination rates, and entity/relation resolution.
- Designed pipelines for aggregating various NLP datasets for specific KGC paradigms, such as Wikidata, Webnlg, and NYT.

TEACHING EXPERIENCE

Brown University, Providence, RI
- *Teaching Assistant for DATA2060: Machine Learning* — *Sep 2025 – Dec 2025*
- *Teaching Assistant for CSCI1460: Computational Linguistics* — *Jan 2026 – May 2026*

ACADEMIC PROJECTS

Soft Token Learning — *Sep 2025 – Dec 2025*
- Investigated soft token learning as a parameter-efficient mechanism for model adaptation, targeting knowledge injection without full fine-tuning or repeated retrieval.
- Proposed training entity-specific soft tokens using structured supervision from knowledge graph paths, encoding contextual semantics directly at the embedding level.
- Analyzed generalization, controllability, and interference behavior, framing the approach as a neural–symbolic interface between structured knowledge and pretrained language models.

Hopfield Networks — *Feb 2025 – May 2025*
- Led a team of three to implement Modern Hopfield Networks (dense associative memory) in TensorFlow, adapting the PyTorch framework from "Hopfield Networks Is All You Need".
- Studied classical Hopfield networks through the lens of energy landscapes and convergence dynamics, and connected these principles to modern Hopfield architectures and transformer attention mechanisms.
- Evaluated memory capacity, retrieval accuracy, and convergence behavior on CIFAR-10 for class-conditional image

reconstruction.

Image Captioning *Feb 2025 – Apr 2025*

- Developed an image captioning system on the Flickr8k dataset using a customized TensorFlow framework.
- Implemented an encoder–decoder architecture with ResNet-50 image encoders and RNN- or Transformer-based decoders.
- Compared RNN and Transformer decoders quantitatively and visualized attention mechanisms to analyze model behavior.

Seq2Seq Semantic Parsing *Nov 2024 – Dec 2024*

- Replicated the neural attention–based semantic parsing model from "Language to Logical Form with Neural Attention," mapping natural language utterances to executable logical forms.
- Implemented a multi-layer LSTM encoder–decoder architecture with attention applied at the decoder.
- Evaluated on the Jobs dataset, achieving over 80% exact-match accuracy within 20 training epochs.

BERT Fine-tuning *Oct 2024 – Nov 2024*

- Fine-tuned a pretrained BERT model with partially frozen parameters for sentiment classification task on tweets.
- Evaluated performance against bag-of-words baselines, training-from-scratch BERT models, and full fine-tuning approaches.
- Demonstrated that parameter-efficient fine-tuning achieves comparable accuracy to full fine-tuning while reducing computation by over 40%.

Classic Information Retrieval System *Jan 2024 – Apr 2024*

- Implemented an ad hoc information retrieval system using TF–IDF weighting and cosine similarity ranking.
- Applied NLP preprocessing with NLTK, including tokenization and stop-word filtering, to improve retrieval accuracy.
- Evaluated on the Cranfield Collection, achieving a mean average precision (MAP) over 30%.

Intelligent Object Detection Rover *Oct 2023 – Dec 2023*

- Implemented a Python-based object detection pipeline for real-time rover perception, enabling autonomous navigation and trajectory following using the Viam platform.
- Integrated hardware and software components, including a webcam, a LiDAR sensor, and the rover platform, into a unified autonomous system.

KTV Web Server *Oct 2023 – Dec 2023*

- Developed a Next.js–based web application for searching and streaming KTV song videos.
- Implemented stateless user authentication and session management using JSON Web Tokens.
- Integrated the YouTube API for video retrieval and MongoDB Atlas with Socket.io to support real-time, multi-user room interactions.

SKILLS

Programming Languages: Python, C++, Java, JavaScript, Go, Scala, SQL, Shell scripting
Machine Learning Frameworks: PyTorch, TensorFlow, Scikit-learn, LangChain, Transformers
NLP: Hugging Face Datasets, Tokenizers, NLTK, spaCy
Web Development: Node.js, React.js, Next.js, HTML, CSS, Tailwind
Databases & Distributed Computing: MongoDB, PostgreSQL, Neo4j, Hive, Spark
Tools & Platforms: Git, Linux, Docker
Language: English, Chinese