

Yuandong Zhang

+1 (858) 214-6408 | ✉ yuz371@ucsd.edu | 🌐 yuandong.me

EDUCATION

University of California San Diego
M.S. in Computer Science, GPA: 4.0/4.0

Expected June 2027
San Diego, USA

Duke University & Duke Kunshan University
B.S. in Computer Science, GPA: 3.8/4.0, Cum Laude

May 2025
Durham, USA

SKILLS

Cloud & DevOps: Kubernetes, Docker, ArgoCD, Amazon EC2, CI/CD Pipelines, GitOps, Linux
Data & AI: Dask, Apache Spark, MongoDB, PyTorch, Scikit-learn, LLMs, FastAPI, OpenCV
Languages: Python, Java, JavaScript/TypeScript, Go, C, SQL, Swift

SELECTED PROJECTS

Enterprise ERP Infrastructure & SaaS | Weifang Royalroad Trading Co., Ltd. May 2025 - Sept 2025

- Built a custom ERP for an export trading company, automating quotation, invoicing, and factory order lifecycle tracking; implemented CRM modules and a historical price analytics engine to optimize cross-client margin strategy.
- Designed and deployed an ETL pipeline to ingest and normalize legacy Excel/PDF transaction records into MongoDB, including schema redesign and data validation, enabling reliable historical pricing analytics.
- Optimized MongoDB aggregation pipelines and indexes for large-scale transaction data, reducing query latency to <100ms.

Smart Cafeteria IoT Recognition System | Yidian Information Technology Co. Ltd. May 2024 - Sept 2024

- Built an end-to-end AIoT checkout system, integrating camera-equipped smart scales with a cloud-based inference backend
- Developed a CNN classifier optimized with Bag-of-Visual-Words (BoVW) to enhance dish recognition accuracy
- Engineered real-time pipeline: edge image capture, cloud inference, and API callbacks for pricing in <100ms

RESEARCH EXPERIENCE

Transformer for Human Mobility Prediction | Duke Kunshan University July 2024 – May 2025

- Engineered a distributed ETL pipeline using Dask to process 4TB of raw GPS trajectories into features in <20 minutes
- Designed a Transformer network achieving 96% SOTA accuracy on trip mode detection using only raw speed inputs
- Benchmarked system performance against ViT, LSTM, and CNN baselines, achieving superior accuracy on Geolife datasets
- Deployed the inference pipeline to handle cross-regional data, boosting generalization accuracy from 80.5% to 86.1%

Federated Black-box Prompt Tuning for LLM | Duke Athena AI Institute Sept 2023 – May 2024

- Designed a Federated Learning system for black-box LLMs, enabling prompt optimization across heterogeneous architectures
- Implemented prompt-projection networks to address data privacy and communication constraints in edge computing settings
- Validated cross-model prompt transfer from BERT/GPT-style models to LLaMA on standard NLP benchmarks
- Reduced communication overhead significantly compared to gradient-based methods by transmitting only prompt vectors

Carbon-Footprint Tracking App | Duke Kunshan University Oct 2023 – Aug 2024

- Engineered a real-time AI pipeline that ingests raw speed data, queries the model, and predicts trip mode instantly
- Deployed the Transformers model using FastAPI and Docker, optimizing serving to handle concurrent API requests
- Led a 7-person team to build the React native app, enforcing strict data privacy compliance for US and China users
- Implemented an item-based recommendation engine to personalize content, driving in-app engagement up 50%

Automated Climate Misinformation Analysis | Duke Nicholas School May 2023 – Jan 2024

- Built a fault-tolerant data pipeline to continuously scrape Twitter streams, bypassing limits via automated queuing
- Deployed a BERT classifier to classify 1.2M+ tweets in real-time, analyzing evolving misinformation trends with precision
- Fine-tuned GPT-3 to separate climate believers from deniers, identifying politicized clusters effectively

PUBLICATIONS & MANUSCRIPTS

- Ding Ma, **Yuandong Zhang**, and Ji Nie. Amplifying Variability of the Southern Annular Mode in the Past and Future. *Geophysical Research Letters*, 2026. <https://doi.org/10.1029/2025GL119214>
- **Yuandong Zhang**, et al. Detecting Transportation Mode Using Dense Smartphone GPS Trajectories and Transformer Models. *International Journal of Geographical Information Science*, 2026. (Accepted)
- Yiming Li, Jingwei Sun, Yudong Liu, **Yuandong Zhang**, et al. Federated Black-box Prompt Tuning System for Large Language Models on the Edge. *MobiCom 2024*