Wenshuo Wang

202364870251@mail.scut.edu.cn https://vincewangvw.github.io

EDUCATION

South China University of Technology, Guangzhou, China **2023.9 - Present**

B.Eng, School of Future Technology, Major in Artificial Intelligence

- **GPA:** 3.6/4.0
- **Major Courses:** Engineering Math: Calculus II, Engineering Math:Probability & Mathematical Statistics, Discrete Mathematics, Introduction to Programming Using Python, Data Structures, Machine Learning, Deep Learning and Computer Vision, Fluid Mechanics, Analysis and Applications of Heat, and related subjects
- **Languages:** English (TOEFL iBT 101)

Massachusetts Institute of Technology, Cambridge, MA, USA **2025.3 - Present**

Visiting Research Intern (Remote), Kavli Institute for Astrophysics and Space Research (Advisor: Dr. Fan Zhang)

RESEARCH INTERESTS

My primary research goal is to **advance efficient scientific computing through deep learning**. I am particularly interested in: **Spatiotemporal Forecasting** (using deep learning models to iteratively predict spatial data at subsequent time steps, such as flow fields, weather, or traffic flow), **Scientific Representation Learning** (finding suitable representation forms for physical field data that are appropriate for model learning), and **Vision-Language Models** (for physical field simulation).

PUBLICATIONS AND PREPRINTS

[1] Breaking Scale Anchoring: Frequency Representation Learning for Accurate High-Resolution Inference from Low-Resolution Training

ICLR 2026 Submission (Under Review) [Paper]

Wenshuo Wang, *Fan Zhang*[†] *([†] Corresponding Author)*

- Reveals that when deep learning models perform spatiotemporal forecasting at higher resolutions, they do not exhibit the power-law error decay characteristic of numerical solvers. Instead, their accuracy remains almost unchanged. We regard this long-accepted behavior as a problem, which we term Scale Anchoring.
- Attribute Scale Anchoring primarily to the Nyquist frequency limit of the training data and introduce Frequency Representation Learning to induce frequency-domain extrapolation with only minor computational overhead, enabling models to mitigate this issue.

[2] Pi-CCA: Prompt-Invariant CCA Certificates for Replay-Free Vision–Language Continual Learning

ICLR 2026 Submission (Under Review) [Paper]

Jiayu Zhang, Chuangxin Zhao, Ruibo Duan, Canran Xiao, Wenyi Mo, Haoyu Gao, ***Wenshuo Wang***[†]

- Diagnose replay-free continual learning of vision–language models as suffering from naive fine-tuning that drifts the image–text alignment geometry and degrades zero-shot transfer and prompt robustness.
- Recast forgetting as alignment-geometry drift and propose PI-CCA, a compact prompt-invariant CCA certificate that preserves canonical spectra and subspaces during adaptation. PI-CCA achieves state-of-the-art replay-free performance with strong zero-shot retention across VL–CL benchmarks.

[3] S²-KD: Semantic-Spectral Knowledge Distillation Spatiotemporal Forecasting

AAAI 2026 [Paper]

Wenshuo Wang, *Yaomin Shen, Yingjie Tan, Yihao Chen*[†]

- Diagnose in existing KD for spatiotemporal forecasting, operate in a semantic vacuum: they can mimic spectral patterns but not the causal "why". This fundamentally limits lightweight students in non-stationary dynamics.
- Introduce S^2-KD: a text-informed multimodal teacher whose fused semantic–spectral representation is distilled into lightweight vision-only students. With simple joint feature and spectral alignment, tiny models nearly match large domain-specialist teachers on WeatherBench, TaxiBJ+ and Prometheus.

[4] Alleviating Choice Supportive Bias in LLM with Reasoning Dependency Generation

Preprint [**Paper**]

Nan Zhuang, **Wenshuo Wang***, Lekai Qian*, Yuxiao Wang, Boyu Cao, Qi Liu† (* Equal Contribution)*

- Diagnose Large language models used as evaluators exhibit strong choice-supportive bias, which the paper reframes as misaligned reasoning dependency on "chosen vs. unchosen" status instead of on evidence.
- Introduce RDG to generate small, structured datasets that reshape these dependencies and fine-tune the model. With only a few thousand RDG examples, choice-supportive bias is sharply reduced while broader social-bias and basic knowledge benchmark metrics remain essentially unchanged.

[5] iTAG: Inverse Design for Text Generation with Accurate Causal Graph Annotations

Preprint [**Paper**]

***Wenshuo Wang**, Boyu Cao, Nan Zhuang, Wei Li†*

- Highlight that real-world text datasets with causal structure annotations are scarce and costly. This fundamentally impedes the advancement of causal discovery in text.
- Propose iTAG, a pipeline that introduces inverse design into text generation. iTAG can transform target directed acyclic graphs into indistinguishable natural text. The synthetic datasets, when applied to existing text causal discovery models, maintain causal discovery metrics that are highly consistent with those on real-world text.

[6] Enhancing Multimodal Emotion Analysis through Fusion with EMT Model Based on BBFN

TCSISR 2024 [**Paper**]

Wenshuo Wang

- Tackle the challenge of accurate multimodal sentiment analysis by fusing text, visual, and acoustic features with an Efficient Multimodal Transformer–enhanced bi-bimodal fusion network to capture cross-modal interactions.

[7] Temporal Data-Driven Short-Term Traffic Prediction: Application and Analysis of LSTM Model

TNS 2023 [**Paper**]

Xuange Huang, Xinran Li*, **Wenshuo Wang****

- Demonstrate that short-term urban traffic flow can be accurately forecast from time-series data alone by training LSTM networks on fine-grained loop-detector measurements and showing they outperform ARIMA baselines.

INTERNSHIP EXPERIENCE

HUAWEI Beijing Research Institute, Beijing, China **2024.07-2024.10**

Algorithm Engineer Intern

- Developed NLP models for user review analysis and medical text understanding, including text vectorization with XGBoost, BERT-based multi-label classification, T5-based summarization, and evaluation of generative models.
- Designed and implemented a medical conversational agent using Streamlit, SimCSE + Faiss retrieval, Alibaba Cloud text embeddings, and GLM-4-0520 for domain-specific question answering.

COMPETITIONS AND AWARDS

- **Kaggle, Yale/UNC-CH - Geophysical Waveform Inversion.** *Bronze Medal, 81/1365, Top 5.9%* **(2025)**
- **Kaggle, NeurIPS - Ariel Data Challenge.** *Silver Medal, 49/1151, Top 4.3%* **(2024)**
- **Kaggle, Home Credit - Credit Risk Model Stability.** *Silver Medal, 23/3856, Top 0.6%* **(2024)**