Xinyue Wang

xiw159@ucsd.edu · xinyuewangg.com · github.com/CharonWangg · Google Scholar · 215.921.1077

Scalable causal learning · World models · Self-improving agents

Education

2023 – 2028 **Ph.D. in Data Science**, UC San Diego (expected early 2028). Advisor: Biwei Huang

2021 – 2023 **M.S.E. in Bioengineering**, University of Pennsylvania. Advisor: Konrad Kording

2017 – 2021 **B.Eng. in Biomedical Engineering**, Shenzhen University; exchange at UPenn (2020)

Experience

Jun – Sep 2026 **Research Scientist Intern**, Aether AI

- Led **TimeBraid**, a family of unified time-series–language models (1.2B–6.7B) for understanding and forecasting, which fuse Qwen3 language towers with TimesFM-derived time-series towers; trained on 2.2M series–text alignment pairs and 4.9M instruction samples, and open-sourced code and 2.5B weights [1].
- TimeBraid-6.7B outperforms every open-weight time-series foundation model (TimesFM-3, Moirai-2, Chronos-2, Toto-2.0) and Claude Opus 4.8 on the GIFT-Eval-CTX forecasting leaderboard (as of Oct 2026), and beats GPT-5.4 by 17.6 points on TSAQA.
- Co-led **RSIAgent**, a recursively self-improving multi-agent system whose curriculum, actor and verifier agents explore new digital environments and distill verified skills into memory, with no training. It lifts open-weight models past GPT-6 Astra on OSWorld 2.0 (79.0% vs. 72.6%); 580+ GitHub stars [2].

Jun – Sep 2025 **Research Scientist Intern**, Abel AI

- Discovered that autoregressive transformers are inherently causal learners: their gradient attributions recover the underlying causal graph with no causal objective (NeurIPS 2025 CauScien workshop) [8].
- Drove multimodal time-series causal discovery over 100K variables for financial insight mining.

Oct 2023 – present **Graduate Researcher**, Causal Intelligence Lab, UC San Diego

- Proposed **WM3C** (ICLR 2025), a world-model framework that lets RL agents generalize to unseen environments by composing language-controlled causal components, with identifiability guarantees; demonstrated in simulation and on real-robot manipulation [4].
- Led **Causal-Copilot**, an autonomous agent that automates end-to-end causal analysis on tabular and time-series data, from data diagnosis and choosing among 20+ methods to self-verification and reporting; 200+ GitHub stars and a live demo at causalcopilot.com [3].

Nov 2021 – May 2023 **Student Researcher**, Kording Lab, University of Pennsylvania

- Proposed learning causal discovery from data instead of hand-designing it: an algorithm trained on interventions in a simulated MOS 6502 microprocessor (3,510 nodes, 1,904 connections) beats human-designed methods on silicon, simulated fMRI and gene networks (TMLR 2023) [9].

2019 – 2021 **Undergraduate Researcher**, MIND Lab, Shenzhen University

- Built a millisecond-level, phase-locked EEG neurofeedback system in C++ on OpenBCI, with 55.6% stronger alpha-wave modulation than prior work (ICBBS 2020).

Selected Publications

Foundation models & LLM agents

[1] **X Wang**, J Pang, K Zhou, K Zhang, D Cao, F Feng, Faisal, S Jin, Y Liu, B Huang. TimeBraid: Unifying Time Series and Language for Understanding and Forecasting. *NeurIPS 2026 Workshop on Foundation Models for Temporal Systems*. [code] [weights]

[2] S Zhu*, S Fan*, **X Wang***, W Wu, K Zhou, B Huang. RSIAgent: Autonomous Exploration for Recursive Self-improvement in New Environments. *Under review*. [code] [project]

[3] **X Wang***, K Zhou*, W Wu*, HS Singh, F Nan, S Jin, A Philip, S Patnaik, H Zhu, S Singh, P Prashant, Q Shen, B Huang. Causal-Copilot: An Autonomous Causal Analysis Agent. *Under review*. [code] [demo]

World models & reinforcement learning

[4] **X Wang**, B Huang. Modeling Unseen Environments with Language-guided Composable Causal Components in Reinforcement Learning. *ICLR 2025*.

[5] Y Yang, B Huang, F Feng, **X Wang**, S Tu, L Xu. Towards Generalizable Reinforcement Learning via Causality-Guided Self-Adaptive Representations. *ICLR 2025*.

[6] Y Wei, K Zhou, L Mao, Z Zhang, Z Xu, Z Xi, S Liang, R Han, Y Yan, **X Wang**, F Feng, B Huang. Causally Debiased Latent Action Model for Embodied Action-Conditioned World Models. *Under review*.

[7] H Liu, F Feng, M Fu, **X Wang**, H Lu, B Huang. SCAR: Self-Supervised Continuous Action Representation Learning. *Under review*.

Scalable causal learning

[8] **X Wang**, S Wang, B Huang. Transformer Is Inherently a Causal Learner. *NeurIPS 2025 Workshop on CauScien*.

[9] **X Wang**, K Kording. Learning Domain-Specific Causal Discovery from Time Series. *TMLR 2023*. [code]

* equal contribution. Full list on Google Scholar.

Skills

Modeling — Unified model learning (time-series–language alignment and instruction tuning, 1.2B–6.7B), LLM agents (multi-agent, self-verification, computer use), world models, reinforcement learning, causal discovery and inference, computer vision

Tools — Python, PyTorch, Hugging Face, C++

Competitions

2021 – 2022 — **Kaggle Expert**, three silver medals: Sartorius Cell Instance Segmentation (17 / 1,559, top 1%), UW-Madison GI Tract Segmentation (32 / 1,548, top 2%), Google chaii Hindi & Tamil QA (39 / 959, top 4%)

Teaching & Service

Teaching — TA for DSC 291 Causal Discovery & Representation Learning (UCSD, 2025; advised graduate research projects), DSC 240 Causal Inference (UCSD, 2024), CIS 522 Deep Learning (UPenn, 2023; mentored 15 students), Neuromatch Academy Deep Learning (2022)

Reviewing — NeurIPS, ICML, ICLR, AISTATS, UAI, CLeaR, RLC, TMLR