

Cheng Xin

AI/ML Engineer · Agentic AI Engineer · Model Evaluation

xin.job2025@gmail.com [jackal092927.github.io](https://github.com/jackal092927) [Google Scholar](https://scholar.google.com/citations?user=...) [GitHub](https://github.com)

SUMMARY

Ph.D. in Computer Science from Purdue University. Currently a **Postdoctoral Researcher at Rutgers University** and **Co-founder/CTO of Synapse Foundry AI**, with **7+ years of AI research and engineering experience** across **agentic AI systems, ML model training and evaluation, large-scale data engineering, graph/topological machine learning, and trustworthy AI**. Publications appeared at **NeurIPS, ICML, CVPR, SoCG** and other top venues; industrial R&D experience at **Amazon, Electronic Arts, and Microsoft**. Led development of **agent gateways, tool-calling and model-service layers, RAG/agentic search, evaluation benchmarks, and auditable workflow automation**. Strong fit for AI/ML engineering, agentic AI engineering, and forward-deployed AI roles that require translating ambiguous customer workflows into data pipelines, model services, evaluation loops, production integrations, observability, and human-in-the-loop handoff systems.

TECHNICAL SKILLS

AI/ML	Predictive Modeling	Decision Modeling	Personalization Modeling	User Behavior Modeling	Graph Learning	Topological ML	Transformer / LLM Evaluation						
Algorithms	Vision-Language Models	Retrieval / Ranking	Non-Euclidean Representation	Explainable AI	Confidence Estimation	Anomaly Detection	Feature Engineering						
LLM Agents	LLM	RAG	Agentic Search	A2A	MCP / Skill	Tool Calling	Context Engineering	Multi-agent Collaboration	Workflow Automation	Human-in-the-loop Handoff			
	Domain Knowledge Bases	Evidence Attribution	Auditable Handoff	Observability									
Development	Python	TypeScript	Node.js	React	SQL	PyTorch	TensorFlow	Spark	AWS	PostgreSQL	API / SDK	Data Pipelines	Automated Evaluation
Tools	Model Deployment	Model Monitoring	MLOps / LLMOps	Customer-facing Prototypes	Full-stack Development	CI/CD							

OPEN SOURCE CONTRIBUTIONS

OpenClaw - Agent Gateway / Multi-agent Infrastructure Open Source Contribution 2026.02 - Present

- As an **early OpenClaw Contributor**, contributed and maintained **30+ Pull Requests**, improving **Agent Gateway, A2A session protocol, message routing, memory-module optimization, model compatibility, and tool calling**; covered interface adaptation, protocol interaction, exception handling, and regression validation in agent infrastructure
- Designed reusable agent skills and automated workflows for task decomposition, cross-source retrieval, long-running task recovery, context handoff, and observability, strengthening runtime patterns for customer-facing agentic applications and long-running AI engineering workflows

DL3DV-10K - Large-scale 3D Vision Dataset and Vision Model Benchmark 2023 - Present

- As a core contributor, helped build a large-scale real-world-scene **3D vision dataset and benchmark** covering **10K+ scenes** and **51M+ video frames**; contributed to data curation, benchmark construction, and experimental validation for large-scale model training and evaluation (**[CVPR 2024]**)
- Project has **600+ GitHub Stars** and has been adopted or cited by industrial and academic teams including **NVIDIA, Adobe Research, Google DeepMind, Meta AI, Microsoft Research, ByteDance, and Tencent Hunyuan**, demonstrating experience with benchmark-driven model evaluation beyond a single research prototype

PeonPing - AI Coding Agent Notification System and Developer Tool Open Source Contribution 2026.02 - Present

- Contributed to a **4.8k Stars** open-source project for mainstream Coding Agent developer tools such as Claude Code, Codex, and Cursor, providing voice/desktop notifications for task completion, permission waits, error states, and session events, helping developers maintain observability and responsiveness during long-running agent tasks
- Implemented configurable notification positions, auto-dismiss/persistent notifications, project label overrides, and message templates; also added CLI subcommands, default configuration, README and bilingual docs, bash/fish completions, and BATS tests, improving configurability, maintainability, and user experience across multi-project, multi-terminal, and long-running agent workflows

INDUSTRY AND FORWARD-DEPLOYED ENGINEERING EXPERIENCE

Synapse Foundry AI - Co-founder / CTO, Agentic AI Systems and Data Gateway Platform 2026.01 - Present

Core technical owner at startup · Led product architecture, R&D implementation, infrastructure buildout, evaluation platform, customer workflow analysis, and delivery-oriented prototypes

- Data/Trust Gateway:** Built a data boundary, evaluation, and governance framework for **agent workflows**, converting customer scenarios, failure cases, permission boundaries, and operating requirements into repeatable automated evaluation systems; constructed a **10k+** benchmark for agent data security and privacy and converted results into coverage, error-distribution, evidence-chain, and issue-localization reports
- Model service and tool-calling layer:** Built an agent tool layer by packaging search, crawl, privacy guardrails, and external capabilities into unified SDK/API interfaces; provided response schemas, cost/latency/quality policies, service routing, failover, call-reason records, and audit logs, forming a reusable model-service and tool-integration layer
- Business workflow automation agents:** Extended large models from generic Q&A to auditable workflow automation systems that can call tools, generate evidence, hand off to humans, and support **raw materials -> information filtering -> tool execution -> evidence generation -> rule checking -> handoff package -> audit summary** closed-loop workflows for domain knowledge Q&A, anomaly handling, report generation, and human-AI decision support
- Clinical Arena (live: clinicalarena.ai):** Built a domain AI evaluation platform with blind head-to-head LLM comparison, preference data collection, safety annotation, leaderboards, and evaluation export; includes RAG-based evidence verification and multi-agent, multi-model collaborative checking to compare answer quality, safety, and stability on real business problems

Electronic Arts: Big Data Intelligence Group - Machine Learning Scientist Intern 2018.05 - 2018.08

- Built an end-to-end **ML pipeline** for player behavior data on **Spark**, covering multi-table extraction, data cleaning, user-profile feature construction, training-sample generation, model training support, and engagement prediction analysis, enabling continuous business-side evaluation of large-scale user behavior
- Used **graph relational models** to analyze connections among users, behavior events, content assets, and relational database tables, restructuring and optimizing user-profile data tables and feature views to improve downstream use of relational data and graph-structured signals
- Decomposed business problems into verifiable data and model tasks, diagnosing sample quality, feature coverage, anomalous distributions, and prediction performance, forming a loop from data preparation and experimental analysis to business interpretation for production ML iteration
- Designed and delivered **feature engineering and data compression** solutions, reducing data-view volume by **40%** and improving downstream analysis efficiency and modeling resource utilization

Amazon: AWS Infrastructure Group - Software Development Engineer Intern

2015.05 - 2015.08

- Developed and deployed a **real-time data management system** for large-scale network message streams, implementing backend infrastructure, service interfaces, message parsing, storage/retrieval, and data-flow modules to support stable internal operations and improve message-processing efficiency
- Optimized interface protocols and data exchange flows for message ingestion, parsing, persistence, querying, and replay around AWS infrastructure data management and data interaction needs, improving consistency, traceability, and maintainability across modules
- Performed debugging and performance optimization for high-throughput message-processing scenarios, focusing on data structures, concurrency, exception recovery, and runtime stability, building engineering foundations for reliable data gateways, model-service interfaces, and scenario toolchains

Koal Software: Security Audit Platform Group - Full-stack Software Engineer

2013.07 - 2013.12

- Owned major Web application modules as a full-stack engineer, covering **requirement decomposition, database modeling, backend business logic, REST-style APIs, frontend pages and interactions, testing/validation, and production delivery**, with end-to-end experience from business requirements to runnable systems
- Implemented Chinese NLP functionality around business text data, including **Chinese word segmentation, keyword extraction, topic modeling, and text-topic understanding**, converting unstructured Chinese text into queryable, analyzable data objects for business judgment
- Connected databases, backend services, NLP processing modules, and frontend presentation layers so that text-analysis results could enter business query, statistical analysis, and user-interaction workflows; this experience connects naturally to current LLM/RAG, domain knowledge base, information extraction, and workflow automation work

Microsoft: SQL Server Business Intelligence Group - R&D Intern

2012.07 - 2012.11

- Served as an R&D intern in the **SQL Server Business Intelligence Group**, working on SQL Server issue diagnosis and feature validation around enterprise data analytics, reporting, query services, and database downstream applications, covering core database objects and mechanisms such as schemas, indexes, views, query plans, logs, and configurations
- Analyzed real enterprise customer and developer issues including slow queries, index failures, view/aggregation logic, data access anomalies, and performance bottlenecks; used execution plans, logs, and environment configurations to identify root causes and propose reproducible, verifiable fixes
- Built systematic understanding of relational database internals and business-intelligence application chains, covering data modeling, query optimization, index/view design, performance diagnosis, and database-driven downstream business application delivery

RESEARCH AND SCENARIO MODEL R&D EXPERIENCE

Rutgers University, Department of Computer Science - Postdoctoral Researcher

2023.10 - Present

Advisor: Prof. Jie Gao

- Developed prototype experiments and evaluation studies around **Agentic RL / multi-step decision-making**, focusing on feedback-signal design, behavior optimization, and experimental analysis; related cooperative learning work for networked agents maps to multi-actor decision-making, scheduling optimization, and human-AI collaboration in industrial scenarios, supporting task decomposition, constraint modeling, and feedback-loop design ([**ICML 2024**])
- Led development of **TopInG**, an interpretable graph learning framework, achieving up to **20%** improvement in accuracy and interpretability on molecular property prediction; for **AI for Science** and biomedical scenarios, applied **GNN** and interpretability methods to real medical problems, collaborating with **Harvard Medical School** on method design, experimental validation, and deployment-oriented application scenarios; transferable to equipment relation networks, process knowledge graphs, root-cause analysis, and expert-reviewable predictive models, while supporting model evaluation and stability analysis on relational data ([**ICML 2025**])
- Contributed to **DL3DV-10K**, a large-scale 3D vision dataset and evaluation benchmark, supporting real-scene data curation, benchmark construction, and experimental validation, with results published at **CVPR 2024**; the dataset has been used by **NVIDIA, Adobe, Google**, and other teams in multiple commercial vision-model and spatial-intelligence scenarios, connecting 3D vision, video understanding, and vision model evaluation
- Conducted research on **non-Euclidean** and **hyperbolic representation learning**, advancing theoretically grounded representation learning and nearest-neighbor retrieval methods around **Neuc-MDS**, Johnson-Lindenstrauss extensions beyond Euclidean geometry, and **Hyperbolic Space LSH**; applicable to complex hierarchical data, knowledge structures, high-dimensional retrieval, similar-case recall, and evidence localization ([**NeurIPS 2024**], [**NeurIPS 2025**], [**SoCG 2026**])

GraphEvol / Complex System State Modeling - Research and Engineering Project

2026 - Present

- Built a graph evolution learning and stage-recovery experimental pipeline that organizes graph-structured data, topological/graph features, model training, and confidence analysis into reusable workflows for modeling state transitions in complex systems, forming an evaluation loop from data ingestion and feature construction to training validation and result analysis
- Designed evaluation methods for structural perturbation, stage-recovery stability, and confidence outputs so model results are not only predictive but also diagnosable and reviewable; transferable to equipment state prediction, operating-condition evolution recognition, anomaly-stage detection, and complex process decision support in industrial settings with multi-source sensor data and stage-based state judgment

Purdue University, Department of Computer Science - Ph.D. Research Assistant

2020.08 - 2023.08

Advisor: Prof. Tamal K. Dey

- Led research on the **GRIL** topological vectorization framework, completing the loop from theoretical modeling and algorithm design to experimental validation and proving its stronger expressive power; extended the work toward **D-GRIL** end-to-end topological learning, enabling multiparameter persistence representations to enter differentiable learning workflows and strengthening graph/geometric models for scientific data, material structures, and complex-system data; applicable to feature engineering and model generalization analysis for complex relational data ([**PMLR 2023**], [**SoCG 2026**])
- Developed a **generalized persistence algorithm** that improves computational efficiency for multiparameter topological analysis tasks, providing reusable algorithmic foundations for large-scale scientific/geometric data processing, model feature construction, and explainable analysis, and serving as a bottom-layer method reserve for complex structured data modeling, compressed representation, and reviewable analysis ([**JACT 2022**])

EDUCATION

Ph.D. in Computer Science · Purdue University

2023

Dissertation: Decomposition and Stability of Multiparameter Persistence Modules · Advisor: Prof. Tamal K. Dey

M.S. in Computer Science · Lehigh University

2016

Thesis: Machine Learning Techniques for Medical Image Analysis · Focus: Computer vision and deep learning for medical imaging

B.E. in Software Engineering · Tongji University

2013

SELECTED PUBLICATIONS

SoCG 2026 C. Deng, J. Gao, K. Lu, F. Luo, C. Xin[†]. "Locality Sensitive Hashing in Hyperbolic Space"

SoCG 2026 S. Mukherjee, S. N. Samaga, C. Xin, S. Oudot, T. K. Dey. "D-GRIL: End-to-End Topological Learning with 2-parameter Persistence"

ICML 2025 C. Xin, F. Xu, X. Ding, J. Gao, J. Ding. "TopInG: Topologically Interpretable Graph Learning via Persistent Rationale Filtration"

NeurIPS 2025 C. Deng, J. Gao, K. Lu, F. Luo, C. Xin[†]. "Johnson-Lindenstrauss Lemma Beyond Euclidean Geometry"

NeurIPS 2024 C. Deng, J. Gao, K. Lu, F. Luo, H. Sun, C. Xin[†]. "Neuc-MDS: Non-Euclidean Multidimensional Scaling Through Bilinear Forms"

ICML 2024 S. Haddadan, C. Xin, J. Gao. "Optimally Improving Cooperative Learning in a Social Setting"

CVPR 2024 L. Ling, ..., **C. Xin**, et al. "DL3DV-10K: A Large-Scale Scene Dataset for Deep Learning-Based 3D Vision"

TMLR 2024 S. Zhang, **C. Xin**, T. K. Dey. "Expressive Higher-Order Link Prediction through Hypergraph Symmetry Breaking"

ICML-W 2023 **C. Xin**, S. Mukherjee, S. N. Samaga, T. K. Dey. "GRIL: A 2-parameter Persistence Based Vectorization for Machine Learning"

† indicates alphabetic author ordering following theoretical research conventions, with C. Xin as corresponding author. Full publication list available on Google Scholar.

HONORS AND LEADERSHIP

Shanghai Magnolia Talent Program - Young Talent (2025)

Microsoft Collegiate Programming Competition Champion (2017 @ Ohio State University, 2015 @ Lehigh University)

Graduate Course Instructor: Design and Analysis of Algorithms (45 students, 2025)

Area Chair: TAG-DS Workshop (2026)

Reviewer: ICML, ICLR, NeurIPS, SoCG