Yunqi He

Email : yunqi.he@u.northwestern.edu
https://actasclown.github.io
Mobile : +1-773-290-0801

Education

Northwestern University — **Evanston, IL**
Ph.D. Candidate in Computer Engineering, advised by Prof. Hai Zhou — *Jan. 2022 – Dec. 2026 (expected)*

Northwestern University — **Evanston, IL**
M.S. in Computer Engineering — *Sept. 2019 – Dec. 2021*

Peking University — **Beijing, China**
B.S. in Computer Science — *Sept. 2013 – Jun. 2017*

Research Summary

I build **physical-design methodology** and ML-driven automation to improve SoC **PPA** across placement, legalization, timing, and routing, emphasizing robust, reliable design flows and applying data science / machine learning where it moves power, performance, and area. Selected directions:

- **Placement & legalization** — a TwinTree-based quality-preserving legalizer and GPU-accelerated analytical placement that turn continuous placements into legal layouts while preserving wirelength, area, and timing.
- **ML for physical design** — GNNs for timing-driven placement initialization and for routing-congestion prediction, learning directly from netlist and layout features.
- **Agentic design automation** — an LLM-driven agent that orchestrates ML-driven engines and commercial EDA tools across the floorplan-to-route flow under unified PPA objectives.

Work Experience

Cadence Design Systems | *Research & Development Intern* — *Jun. 2026 – Dec. 2026*

- Building **GNN-based routing congestion prediction** for the Palladium (Z4) emulator: forecasting congestion from netlist and placement features to guide routing and improve routability on industrial designs.
- Returning intern extending prior GPU placement-acceleration work into the **routing** stage of the design flow.

Cadence Design Systems | *Research & Development Intern* — *Jun. 2024 – Sept. 2024*

- Integrated GPU-accelerated placement (DREAMPlace) into Palladium (Z3) design flows, achieving **6–12**× speedup over CPU-based baselines on multi-million-cell industrial designs while maintaining solution quality and convergence stability.
- Analyzed **CUDA kernel** performance and memory behavior to identify optimization opportunities for large-scale mixed-size placement workloads.
- Developed interfaces bridging academic GPU-accelerated kernels with commercial flows, enabling heterogeneous CPU–GPU execution and rapid prototyping of ML-enhanced placement in production-like environments.
- Collaborated with architects and CAD engineers to assess scalability and numerical stability of GPU-accelerated optimization and its integration with existing timing and congestion analysis engines.

Selected Research Projects

Ph.D. Dissertation: An Agentic System for End-to-End VLSI Physical Design — *Expected Dec. 2026*

- Building an **agentic EDA system** in which an LLM-driven agent plans and executes the full physical-design flow — floorplanning, placement, legalization, and routing — invoking ML-driven engines and commercial tools under unified **PPA** objectives.
- Grounded in core contributions the agent invokes as reliable primitives: a **TwinTree-based quality-preserving legalizer** and **GNN-guided** design engines (detailed below).
- Keeping agent actions on **physically valid states** via built-in topological invariants and solver-based checking of constraint violations.

Quality-Preserving Legalization for Analytical Floorplanning [1] — *Jun. 2024 – Apr. 2026*

- Built on a TwinTree representation of floorplan topology that supplies **built-in topological invariants**, with a TAG graph abstraction layered on TwinTree to capture adjacency and geometric relationships among modules.
- Designed a **quality-preserving legalization** procedure on TwinTree that resolves overlaps and boundary violations while staying close to the analytical placer's solution, preserving wirelength and area.
- Integrated the legalizer with analytical floorplanning and GPU-accelerated backends, providing a **physical-aware legalization primitive** validated on industrial-scale benchmarks.

GNN-Guided Physical Design: Timing-Driven Placement & Floorplanning [2] *2025 – Present*

- Graph neural networks that learn cell- and net-level embeddings from the netlist to produce **timing-aware initial placements**, improving downstream timing and convergence of the analytical placer.
- Framed placement initialization as node-level prediction on the netlist graph; the learned seed warm-starts the placer, cutting placement iterations and improving post-placement timing over conventional initialization.
- Extending the GNN line to floorplanning — building a **data-driven, agile ML-based floorplanning tool** that learns from the FloorSet dataset of real-world SoC layouts to guide block placement, providing the dissertation's floorplanning engine paired with the TwinTree legalizer.

ML-Guided Hardware Automation & Secure Redaction Pipelines [3–5] *May 2021 – Sept. 2025*

- Designed GNN models over gate-level netlists and placement features to guide **constrained region selection**, casting redaction as a graph-partitioning problem with embedding-driven clustering under area and timing budgets.
- Implemented a multi-stage ML–EDA flow interleaving GNN-guided decisions with Yosys, OpenSTA, and OpenROAD, showing how strongly EDA quality depends on physical-awareness in ML representations.
- Closed the loop with OpenFPGA-based fabric generation, yielding an end-to-end pipeline for secure, verifiable design transformations that balances area, timing, and security under explicit budgets.

Automated Verification for Hardware Design Transformations [6, 7] *Mar. 2022 – Dec. 2024*

- Built automated **equivalence-checking** flows for sequential, non-cycle-accurate design transformations around commercial tools (Synopsys Hector, Cadence Jasper, Mentor SLEC), using refinement mapping.
- Constructed benchmark suites and experimental harnesses to evaluate correctness, soundness, and scalability across sequential and non-cycle-accurate transformations.

Earlier Hardware & RTL Design (FPGA / Verilog / HLS) *2016 – 2017*

- **FPGA CNN accelerator**: mapped neural-network datapaths to FPGA hardware with Xilinx Vivado HLS, applying array partitioning and loop unrolling to accelerate the generated logic.
- **SDR physical-layer RTL**: extended the Verilog RTL of an FPGA software-defined-radio platform (GRT), exporting channel-state information from the receive chain through a custom serial I/O readout path.

Technical Skills

Scripting & Programming: Python, Tcl, C/C++, Verilog/SystemVerilog/VHDL

Physical Design & PPA: physical-design & production flows, placement & routing (Cadence Innovus, OpenROAD, DREAMPlace), cell legalization, static timing analysis (OpenSTA), congestion analysis, logic synthesis (Cadence Genus, Synopsys DC, Yosys)

Data Science & ML: PyTorch, graph neural networks (GNNs), LLM-based agents applied to PD/PPA; CUDA, GPU acceleration

Verification: equivalence checking & formal methods (Cadence Jasper, Synopsys Hector, Mentor SLEC)

Systems: Linux, Git, Docker; multithreading, GPU clusters

Relevant Coursework

Hardware & VLSI: VLSI Algorithmics, Formal Verification, ASIC & FPGA Design, Reconfigurable Systems

Systems & Programming: Massively Parallel Programming with CUDA, Multicore Concurrent Programming

Machine Learning & AI: Machine Learning, Artificial Intelligence, Probability Theory & Statistics, Random Processes

Teaching & Service

Graduate Teaching Assistant: Design Automation in VLSI, Advanced Digital Design, Computer System Software, Intro to Computer Engineering

Conference Reviewer: ICCAD 2026, ASP-DAC 2026, HOST 2024/2025, AAAI 2026

Awards

- Northwestern University Terminal Year Fellowship (2026)
- "Excellent Graduation Design" (top 5%) of EECS Dept., Peking University (2017)
- Northwestern University Ph.D. Fellowship (2022)
- 3rd Prize, ACM Programming Contest of Peking University (2015)

Publications

[1] **Yunqi He**, You Li, Ruofan Huang, and Hai Zhou. Quality-preserving legalization for analytical floorplanning. In *International Conference on Computer-Aided Design (ICCAD)*. IEEE/ACM, 2026. To appear.

[2] Ziyi Ju, **Yunqi He**, Ping Yu, Hai Zhou, Jia Wang, and Fang Yan. Graph neural network based initialization for timing driven placement. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, Apr. 2026.

[3] **Yunqi He**, You Li, Ruofan Huang, Guannan Zhao, and Hai Zhou. Physical-aware eFPGA redaction for secure and efficient hardware IP protection. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2026.

[4] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. DE2: SAT-based sequential logic decryption with a functional description. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2025.

[5] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. ObfusLock: An efficient obfuscated locking framework for circuit ip protection. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2023.

[6] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. SE3: Sequential equivalence checking for non-cycle-accurate design transformations. In *Design Automation Conference (DAC)*. ACM/IEEE, 2023.

[7] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. RE3: Finding refinement relations with relational mapping abstraction. In *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, pages 1–7. IEEE, 2025.

[8] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. Evaluating the security of logic locking on deep neural networks. In *Design Automation Conference (DAC)*. ACM/IEEE, 2024.

[9] You Li, Guannan Zhao, **Yunqi He**, and Hai Zhou. Certifying global robustness for deep neural networks. *arXiv preprint arXiv:2405.20556*, 2024.

[10] You Li, Guannan Zhao, Yuhao Ju, **Yunqi He**, Jie Gu, and Hai Zhou. LLA: Enhancing security and privacy for generative models with logic-locked accelerators. In *AAAI Conference on Artificial Intelligence*. AAAI, 2025.