Peijing Li

peli@stanford.edu | +1 (515) 715-3989 | linkedin.com/in/peijing-li | peijli.github.io

Education

Stanford University — Stanford, CA
M.S. in Electrical Engineering (expected Dec 2026) — Sep 2024 – Dec 2026
- Concurrently in the Ph.D. program since Sep 2024

University of Michigan — Ann Arbor, MI
B.S.E. in Computer Science — Aug 2020 – May 2024
- Minor in Civil Engineering; Cumulative GPA 4.0/4.0; graduated summa cum laude

Research Experience

Stanford Differentiated Access Memories Project — Sep 2024 – Mar 2026
- **GainSight:** first author of an open-source profiler that measures cycle-accurate data lifetimes in accelerator workloads and generates optimal heterogeneous on-chip memory compositions
- **Flan:** built an MLIR compiler pass that solves heterogeneous memory placement as an Optimization Modulo Theories problem in Z3, deriving migration, refresh, and eviction policies from the solver instead of hand-written heuristics
- **GEMMA:** redesigned the embedded DRAM partition, design-for-test module, and refresh controller for a 16 nm AI accelerator tapeout; implemented MXFP6 arithmetic and authored the AXI and DMA transaction specifications
- **Publications:** six publications, including two peer-reviewed workshop papers; first author on two. Full list at peijli.github.io

Stanford Undergraduate Visiting Research Program (UGVR) — Jun 2025 – Aug 2025
- Mentored a visiting undergraduate researcher through a ten-week project, defining the problem scope and onboarding the student to two active projects
- Guided work that added cycle-accurate memory-access logging to the Synopsys ASIP Designer toolchain by instrumenting the RTL testbench memory module to emit load and store records
- Scoped a DMA-based DRAM subsystem for a VLIW and RISC-V AI core, including a three-instruction custom ISA extension and an HBM-parameterized latency model
- Profiled four processor designs across workloads ranging from single CNN layers up to Llama-3-8B

Industry Experience

TSMC North America — San José, CA
AI Hardware Research Intern — Jun 2026 – Sep 2026
- Extended an in-house analytical performance model for ASIC and GPU inference to support workload-level profiling
- Investigated embedded DRAM opportunities across activation, KV cache, and weight data types for large language model serving
- Built an analytical cost model for multi-chip scaling behavior, unifying several interconnect topology families under a single latency and energy notation

Stanford University Department of Computer Science — Stanford, CA
Teaching Assistant, CS 217 Hardware Accelerators for Machine Learning — Jan 2026 – Mar 2026
- Rebuilt the course with two faculty and three TAs after a 4-year hiatus; enrollment hit 100 against a forecast of 40 to 60
- Authored the SystemC and Verilog labs in which students build an LLM accelerator one component at a time: systolic array, memory hierarchy, compiler backend
- Tuned DSP mapping in Siemens Catapult and Xilinx Vivado to keep FPGA lab turnaround tractable at course scale
- Advised students on final course projects by designing project rubrics and reviewing periodic checkpoint deliverables

ASML Silicon Valley — San Jose, CA
Software Test Engineer Intern — May 2023 – Aug 2023
- Automated regression testing triage for focus-exposure modeling software
- Replaced 2-3 hours of daily manual triage, split across five engineers, with an unattended pipeline
- Attributed each failure in a 300-case nightly suite to its commit and owner in under ten minutes
- Deployed the scripts into the CI/CD pipeline with automated stakeholder reporting

Dell Technologies — Shanghai, China
Software Engineering Intern, Research & Development — Jun 2021 – Aug 2021
- Automated VxRail lifecycle-management validation across the Kubernetes interface, VMware ESXi, and VMware Cloud Foundation using Python and Selenium
- Cut a 3-4 hour manual procedure to roughly 30 minutes running unattended
- Debugged physical server racks when experimental deployments failed to boot

Skills

- **Programming:** Python, C/C++, CUDA, RISC-V/ARM/x86 assembly, PyTorch, JAX
- **Hardware:** SystemVerilog, Verilog, SystemC, high-level synthesis (Catapult, Vivado), Synopsys ASIP Designer
- **Compilers & Tools:** LLVM, MLIR, Z3, SMT, CI/CD, Git, Docker, Selenium, MATLAB