Last updated in November 2026

Junseo Lee

junseo.lee@snu.ac.kr | junseolee0.github.io | Google Scholar

RESEARCH INTERESTS

Computer Architecture, GPU Microarchitecture, Hardware-Software Co-Design for Emerging Workloads, Computer Vision/Graphics, Neural Rendering

EDUCATION

Seoul National University — Seoul, Republic of Korea

M.S./Ph.D., Electrical and Computer Engineering — Mar 2022 – Feb 2028 (expected)

- Computer Architecture and Systems Lab (Advisor: Prof. Jaewoong Sim)

B.S., Electrical and Computer Engineering — Mar 2018 – Feb 2022

PUBLICATIONS

HPCA '27 *to appear*
BlendCore: Architectural Support for High-Throughput Programmable Blending
Junseo Lee, Wonjae Jang, Sangyun Jeon, Jaewoong Sim
33rd International Symposium on High-Performance Computer Architecture, 2027

HPCA '26
GRTX: Efficient Ray Tracing for 3D Gaussian-Based Rendering
Junseo Lee, Sangyun Jeon, Jungi Lee, Junyong Park, Jaewoong Sim
32nd International Symposium on High-Performance Computer Architecture, 2026

HPCA '25
VR-Pipe: Streamlining Hardware Graphics Pipeline for Volume Rendering
Junseo Lee, Jaisung Kim, Junyong Park, Jaewoong Sim
31st International Symposium on High-Performance Computer Architecture, 2025

ASPLOS '24
GSCore: Efficient Radiance Field Rendering via Architectural Support for 3D Gaussian Splatting
Junseo Lee, Seokwon Lee, Jungi Lee, Junyong Park, Jaewoong Sim
2024 International Conference on Architectural Support for Programming Languages and Operating Systems, 2024

ISCA '23
NeuRex: A Case for Neural Rendering Acceleration
Junseo Lee, Kwanseok Choi, Jungi Lee, Seokwon Lee, Joonho Whangbo, Jaewoong Sim
50th Annual International Symposium on Computer Architecture, 2023

HONORS AND AWARDS

Outstanding Graduate Student Research Fellowship: Department of ECE, Seoul National University — 2026

Student Travel Grants: ISCA '23, ASPLOS '24, HPCA '25 — 2023–2025

RESEARCH EXPERIENCE

Graduate Research Assistant Mar 2022 – Present
Seoul National University (Advisor: Prof. Jaewoong Sim) Seoul, Republic of Korea

- **BlendCore (HPCA '27): Extending the graphics pipeline in GPUs for high-throughput shader blending**
  - As a follow-up to VR-Pipe, offloaded blending from the fixed-function ROP to high-throughput SMs to resolve the blending bottleneck.
  - Designed a lightweight tile distribution unit that dynamically steers blending work to lightly loaded SMs with only a few kilobytes of storage per Graphics Processing Cluster (GPC).
- **GRTX (HPCA '26): Software and hardware optimizations for efficient ray tracing of 3D Gaussians**
  - Proposed a streamlined acceleration structure in which all Gaussians share a single unit-sphere geometry with per-Gaussian transformations to substantially reduce BVH size and traversal overhead.
  - Designed hardware support for traversal checkpointing in ray tracing units to eliminate redundant node visits during multi-round tracing.
- **VR-Pipe (HPCA '25): A hardware extension to the graphics pipeline in GPUs for volume rendering**
  - Extended the hardware graphics pipeline to reduce the pressure on the fixed-function blending unit (ROP).
  - Designed hardware support for early termination to skip unnecessary fragments and quad merging to pre-blend fragments in shader cores before the ROP.
- **GSCore (ASPLOS '24): A specialized accelerator for the 3D Gaussian Splatting rendering pipeline**
  - Proposed three algorithmic optimizations and tailored hardware to reduce ineffective computations.
  - Showed that the algorithmic optimizations alone speed up rendering on GPUs in software.
- **NeuRex (ISCA '23): Algorithm-hardware co-design for Neural Radiance Fields (NeRF) rendering**
  - Made the multi-resolution hash encoding hardware-friendly and co-designed an encoding engine, which together eliminate irregular off-chip memory accesses.
  - Implemented NeuRex in SystemVerilog with small area overhead and low power consumption.

TEACHING

Graduate Teaching Assistant
Seoul National University Seoul, Republic of Korea

- ECE 315.A: Digital Systems Design and Experiments Fall 2022
- ECE 322: Computer Organization Spring 2022

Led lab sessions and the final project that implements a CNN accelerator in RTL on an FPGA (ECE 315.A). Prepared assignments, including a C++-based cycle-level cache simulator (ECE 322).

SKILLS

Languages: C/C++, CUDA, Python, Verilog/SystemVerilog, GLSL

Tools: PyTorch, Intel VTune, Nsight Compute/Graphics/Systems

APIs: OpenGL, Vulkan, OptiX