

Ketaki Joshi

ketaki.joshi@yale.edu
+1 203-906-4461
<https://joshi-ketaki.github.io/>

PROFESSIONAL EXPERIENCE

• Nvidia

Unified Virtual Memory Team

June 2022 - August 2022

- Identified the root cause, proposed and delivered a new data eviction policy to improve memory management in tiered memories.
- The improved policy delivers two orders of magnitude performance improvement and enables efficient execution of emerging applications.

Architecture Research Team

June 2021 - October 2021

- Identified and developed a prototype to execute traditional GPU applications on CPU SIMD units to reduce data transfer costs across tiered memory systems without compromising programmability.
- Achieved speedup by a factor of 1.5 compared to pure CPU-SIMD execution and a speedup by a factor of 1.8 compared to a pure GPU execution.

GPU Compiler Team

January 2017 - August 2019

- Delivered compiler frontend and backend interface design, development and code generation for the deep learning matrix operations. The instructions were exposed in CUDA 10.0 and CUDA 10.1. (MMA etc.)
- Delivered the entire assembly code generation and decoding for Turing architecture.
- Led analysis of key benchmarks to identify opportunities within the compiler for a newly introduced architecture feature of uniform register files.
- Designed and deployed a framework to auto-generate assembly and decoding of assembly instructions for the compiler.

Tools Development Team

June 2015 - July 2016

- Identified the root cause and developed a no-reference image analysis tool to identify artifacts in images rendered across a different of GPU architectures.
- Achieved an accuracy of 98% and eliminated the time-consuming manual analysis.
- The work was submitted to Nvidia internal technical conference and an Invention Submission Form was filed.

• Shoreline IoT

Member of Technical Staff

September 2016 - January 2017

- Led the prototype development of a remote monitoring device for industrial boilers.
- Collaborated with the design and development teams to ensure the delivery of a reliable, effective solution which significantly improved the efficiency and safety of boiler operations, leading to substantial cost savings and decreased downtime.

RESEARCH AND TEACHING EXPERIENCE

• Yale University

Post-qualifier Ph.D researcher

August 2019 - Present

- Mitigating Catastrophic Forgetting in Online Machine Learning using Brain-Inspired Techniques.
 - * Developed a technique inspired by cognitive science to reduce catastrophic forgetting in online machine learning applications.
 - * Demonstrated the impact of this strategy for memory access prefetching, a significant issue in computer system optimization.

- * Achieved a speedup of 1.6 times compared to existing regularization techniques to mitigate catastrophic forgetting. Also, reduced external storage by a factor of 16 compared to existing replay techniques in initial prototype evaluations.
- Prefetching using Principles of Cognitive Systems.
 - * Explored the application of brain-inspired, machine learning techniques to enhance memory management in tiered memory systems.
 - * Identified the presence of naturally occurring replay - a catastrophic mitigation tactic - in the GPU memory management model. This addresses the fundamental issue of catastrophic forgetting in online learning, without the need for external implementation.
- Leveraging Computer Systems to Better Understand Cognitive Systems.
 - * Developed a code clone detection tool to identify computationally similar modules.
 - * This enabled neuroscientists to reuse, construct, and comprehend cognitive models.

Teaching Fellow for CPSC 323: Introduction to Systems Programming

- Designed and deployed a toy compiler to introduce students to compiler optimizations and assembly code generation, using a custom-designed toy programming language and a subset of x86 ISA.

• **Indian Institute of Technology, Bombay**

GCC Resource Center

May 2014 - June 2015

- Developed a sophisticated tool that, when provided with data flow equations, can generate optimization passes. These optimization passes are then seamlessly integrated into the compiler. This offers programmability and simplicity, making it easier to experiment with and develop new compiler optimizations.

• **Indian Institute of Tropical Meteorology**

February 2014 - June 2014

- Deployed a sophisticated N-SAT solver to be utilized in weather prediction analysis modules. This solver plays a critical role in enhancing the accuracy and efficiency of these analysis and decision-making modules, improving the overall performance of the weather prediction system.

EDUCATION

• **Yale University**

M.Phil in Computer Science

June 2021 - June 2022

- **Thesis:** “Single Source Code, Hardware Agnostic Heterogeneous Systems.”

Masters in Computer Science

August 2019 - June 2021

- **Thesis:** “Detecting Computational Clones in Brain Models.”

• **University of Pune**

Bachelor's in Computer Engineering

August 2011 - June 2015

Institute Rank : 1/200, University Rank : 5/9000

- **Thesis:** “OptGen : A Custom Compiler Optimization Generator.”

SKILLS, AWARDS, AND SERVICE

• **Domain Tools and Skills:**

- **Programming Languages, Frameworks and Source Control:** C, C++, Python, x86 Assembly, Nvidia PTX Assembly, CUDA, MATLAB, Octave, Pytorch, Raytune, Git, PerForce.
- **Writing and Presentation Tools:** L^AT_EX, MS Office, Libre Office
- **Writing and Presentation Courses:** Academic Writing (Yale Course: CPSC 992), Great Presentations (Yale Course: CPSC 993).

- **Awards :**

- **Ackerman Faculty Award Committee 2022:** Selected to be the sole representative of the entire Computer Science Graduate Community.
- **Travel Grants:** ISCA 2023, ISCA 2022, MICRO 2020, CRA Women Graduates Meet 2020.
- **Professional Recognition:** Recognized as a “Strong Contributor” for work done during full-time employment at Nvidia.
- **Persistent Systems Gold Medal:** Gold Medal for being the 2014 department topper in junior year.

- **Service:**

- **Yale Society for Women Engineers Board:** Currently serving as the Communications Chair.
- **Committee Member on 2022 Ackerman Faculty Award:** Participated as the committee member in evaluating, debating, and finalizing the winning candidate.
- **Nvidia Women in Technology Board Member:** Served on the board in the India chapter.

- **Mentorship:**

- Mentored two women Yale sophomores for navigating internships, career, and academic challenges under the Women in Sciences at Yale (WISAY) mentorship initiative.

PUBLICATIONS AND PRESENTATIONS:

- **Ketaki Joshi**, Raghavendra Pradyumna Pothukuchi, Andre Wibisono, Abhishek Bhattacharjee “Mitigating Catastrophic Forgetting in Long Short-Term Memory Networks.”, arXiv:2305.17244 [cs.LG].
- Wu Michael, **Joshi Ketaki**, Sheinberg Andrew Cox Guilherme, Khandelwal Anurag, Pothukuchi Raghavendra Pradyumna, Bhattacharjee Abhishek, “Prefetching Using Principles of Hippocampal-Neocortical Interaction.”, HOTOS ’23.
- **Ketaki Joshi**, Guilherme Cox, Jan Vesely “Access Guided Eviction for Unified Virtual Memory.”, Summer’22 Internship Talk at Nvidia.
- **Ketaki Joshi**, “Single Source, Hardware Agnostic Heterogenous Systems.”, Fall’21 Candidacy Exam Talk at Yale University.
- **Ketaki Joshi**, Daniel Lustig, Oreste Villa, “CUDA Task launcher for GPU and CPU SIMD units.”, Fall’21 Internship Talk at Nvidia Research.
- J. Vesely, R. P. Pothukuchi, **K. Joshi**, S. Gupta, J. D. Cohen and A. Bhattacharjee, “Distill: Domain-Specific Compilation for Cognitive Models.”, 2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), 2022, pp. 301-312, doi: 10.1109/CGO53902.2022.9741278.
- J. Vesely, R. P. Pothukuchi, **K. Joshi**, S. Gupta, J. D. Cohen, and A. Bhattacharjee, “Cognac: Domain-Specific Compilation for Cognitive Models.”
- **Ketaki Joshi**, Uday Khedker “A Custom Compiler Optimization Pass Generator.”, Fall’ 15 guest speaker at ACM-W chapter of Cummins College of Engineering for Women, Pune University.