

## EDUCATION

<b>Carnegie Mellon University</b>	Pittsburgh, USA	09/2023-present
<i>PhD student, Computer Science Department of School of Computer Science</i>		
<ul style="list-style-type: none"> <li>Advisor: Phillip B. Gibbons, Heather Miller.</li> <li>Expected Graduation: 06/2028</li> </ul>		
<b>Peking University</b>	Beijing, China	09/2019–06/2023
<i>Undergraduate student, School of Electronic Engineering and Computer Science</i>		
<ul style="list-style-type: none"> <li>Major in Computer Science (Turning Class)</li> <li>Graduated Summa Cum Laude</li> </ul>		

## PUBLICATION

(\*Equal Contribution)

- [Siyuan Chen](#), Zhuofeng Wang, Zelong Guan, Yudong Liu, Phillip B. Gibbons. “Practical Offloading for Fine-Tuning LLM on Commodity GPU via Learned Sparse Projectors.” In the proceedings of 39<sup>th</sup> Association for the Advancement of Artificial Intelligence (AAAI-39), February 2025.
- Size Zheng, [Siyuan Chen](#), Siyuan Gao, Liancheng Jia, Guangyu Sun, Runsheng Wang, Yun Liang. “TileFlow: A Framework for Modeling Fusion Dataflow via Tree-based Analysis.” 2023 56th IEEE/ACM International Symposium on Microarchitecture (MICRO). IEEE, 2023.
- [Siyuan Chen](#), Pratik Fegade, Tianqi Chen, Phillip B. Gibbons, Todd C. Mowry. “ED-Batch: Efficient Automatic Batching of Dynamic Deep Neural Networks via Finite State Machine.” International Conference on Machine Learning (ICML). PMLR, 2023.
- Size Zheng, [Siyuan Chen](#), Yun Liang. “COMB: Memory and Computation Coordinated Mapping of DNNs onto Complex Heterogeneous SoC.”, in the proceedings of the Design Automation Conference (DAC-60), July 2023.
- Size Zheng\*, [Siyuan Chen](#)\*, Pedi Song, Renze Chen, Xiuhong Li, Shengen Yan, Dahua Lin, Jingwen Leng, Yun Liang. “Chimera: An Analytical Optimizing Framework for Effective Compute-intensive Operators Fusion”, in Proceedings of the 29th international symposium on High Performance Computer Architecture (HPCA-29), February 2023.

## PREPRINT

- Peter Zhong\*, [Siyuan Chen](#)\*, Ruiqi Wang, McKenna McCall, Ben L. Titzer, Heather Miller, Phillip B. Gibbons. “RTBAS: Defending LLM Agents Against Prompt Injection and Privacy Leakage”
- [Siyuan Chen](#), Zhipeng Jia, Samira Khan, Arvind Krishnamurthy, Phillip B. Gibbons. “SLOs-Serve: Optimized Serving of Multi-SLO LLMs”.

## Professional Experience

<b>Google</b>	New York City, USA	05/2024-08/2024
<i>Student Researcher. System research @ Google (SRG).</i>		
Topic: LLM serving		
Host: Samira Khan.		

## RESEARCH

Interest: Machine Learning System, Large Language Model, Algorithm design

**SLO-Optimized LLM Serving** 05/2023~ 2/2024

*Individual Research, Google and Carnegie Mellon University*

- Heterogeneous LLM applications underscores the need for heterogeneous service level objectives (SLOs), e.g. daily chatting demands generation speed like human’s reading speed, while summarization tasks demand end-to-end latency.
- Customizing SLO for different applications presents unique challenge for scheduling in LLM serving to satisfy request-dependent SLO while sharing resources for maximal throughput.
- Propose a LLM specific scheduling algorithm that performs optimal admission control and perform the resource allocation that handles heterogeneous SLO requirements.
- Achieving upto 2.1x higher request throughput compared to SOTA serving frameworks under SLO requirement.

**Defending LLM Agents Against Prompt Injection and Privacy Leakage**

10/2023~ 2/2024

*Cooperated Research, Carnegie Mellon University*

- Agentic applications allow LLM to use external tools for tasks beyond their standalone capabilities. Tools show severe security concern due to prompt injection attack and accidental privacy leakage.
- Propose to protect agentic applications selective information flow approach that keeps track of information propagation in agentic applications, and block tool calls with low-integrity inputs. Implement the information tracking via attention mechanism.
- Evaluation on SOTA prompt injection benchmark shows the defense successfully blocks all attacks with only 2% loss on utility. Help creating a benchmark on privacy leakage and shows the attention-based information-tracking approach is able to detect privacy leakage with near-human performance.

## Memory efficient LLM fine-tuning on commodity hardware via learned subspace projectors

11/2023~ 6/2024

*Individual research, Carnegie Mellon University*

- Fine-tuning LLMs requires significant memory, often exceeding the capacity of a single GPU. A common solution to this memory challenge is offloading compute and data from the GPU to the CPU. However, this approach is hampered by the limited bandwidth of commodity hardware, which constrains communication between the CPU and GPU.
- We designed an offloading framework, DyServe, that enables near-native speed LLM fine-tuning on commodity hardware through learned subspace projectors.
- As a result, our framework can fine-tune a 1.3 billion parameter model on a 4GB laptop GPU and a 7 billion parameter model on a GPU with 24GB memory. Compared to SOTA offloading frameworks, our approach increases fine-tuning throughput by up to 3.33 times and reduces end-to-end fine-tuning time by 33.1%-62.5% when converging to the same accuracy.
- Accepted to AAAI'25 as first author.

## TileFlow: A Framework for Modeling Fusion Dataflow via Tree-based Analysis

3/2023-7/2023

*Undergraduate Dissertation, Supervised by Prof. Eric Liang, Depart. of EECS, Peking University*

- Observed that though with good empirical performance, fusion dataflow for tensor programs is hard to design and compare
- Formulate and automate the design process of fusion dataflow into three stages: Mem-tree design, Tile-tree design, and Loop-tree design.
- Developed a cycle-level simulator to quantize fusion dataflow designs supporting customized software and hardware description.
- Compare and analyze different dataflow designs on different hardware configurations. Design a new fusion dataflow outperforms SOTA dataflow by 1.3x.
- Awarded the top-10 thesis in School of EECS at Peking University.
- Accepted to MICRO23' as co-author.

## Optimizations for Dynamic Batching algorithm for Dynamic Neural Networks

07/2022-01/2023

*Individual Research, Supervised by Prof. Phillip B. Gibbons, Todd Mowry and Tianqi Chen, Depart. of EECS, CMU*

- Observed current technique to exploit batched parallelism for dynamic neural networks is suboptimal.
- Propose FSM-based dynamic batching to find better batching choice. Automate the FSM discovery by RL-based searching algorithm.
- Extended a data structure PQ-Tree to reduce the memory copy by better memory layout;
- Achieved up to 2.4x end2end speedup for DAG-based dynamic neural networks on CPU and GPU compared to state-of-the-art dynamic frameworks, cut down kernel launches by 30~50%, cut down memory transfer amount by 40~50%;
- Accepted to ICML23' as first author.

## Mapping Heterogeneous Neural Networks onto Heterogeneous SoC

10/2022-11/2022

*Cooperated Research, supervised by Prof. Eric Liang, Depart. of EECS, Peking University*

- Observed current mapping framework does not consider the time sharing within one accelerator and treats communication between accelerators equally.
- Proposed a mapping framework to map heterogeneous neural network onto heterogeneous SoC that considers both DNN operators' time sharing, resource sharing, and SoC's heterogeneity in computation and communication.
- Designed and implemented a genetic algorithm to explore the combinational search space.
- Achieved 1.3~1.5x overall latency gain compared to state-of-the-art mappers.
- Accepted to DAC'23 as co-author.

## Analytical Model on Kernel Fusion for Compute-Intensive Operator Chains in DNN on CPU and GPU

10/2021-6/2022

*Cooperated Research, supervised by Prof. Eric Liang, Depart. of EECS, Peking University*

- Observed that compute-intensive operators (GEMM CONV) become memory intensive in DNN workloads.
- Applied aggressive kernel fusion to compute intensive operator chains for better locality and characterize the design space by an analytical model.
- Designed a constant complexity algorithm to solve for the best loop transformation configuration for fused kernels.
- Implemented an auto-scheduler based on TVM and achieve 1.5~2x speedup on CPU compared to vendor libraries and state-of-the-art tensor compilers.
- Accepted to HPCA' 23 as the common first author.

## Rewards

---

- ML Commons Rising Star 2025.

## Talks

---

- "Towards Securing Tool-usage for Customized LLMs." PDL retreat, Oct. 2024.

## Skills

---

Programming Language: Proficient in C++, python;  
Framework: Pytorch, Tensorflow, TVM.