

Madukuri Shaurya

✉ shaurya.m207@gmail.com ☎ 7075051641 in Madukuri Shaurya

Education

IIT Delhi — *B.Tech in Electrical Engineering, CGPA 7.77*

Nov 2020 – May 2024

Experience

Data Scientist II, *Skan.AI* — Computer Use Agent for BFSI Desktop Automation

March 2026 – Present

- **Leading a team** building a **Computer Use Agent** for automated desktop GUI interaction across **12+** BFSI desktop workflows.
- Architected a **stateless step executor** driven by an Agentic Orchestration Procedure (AOP) — takes a screenshot + step name, runs atomic GUI actions, returns control.
- Implemented **Set-of-Mark (SoM)** prompting with a thinker-grounder model split for visual grounding, reaching **95%** atomic step success; added **HITL escalation** and grace-condition handling for UI collisions and post-action verification.

Data Scientist I, *Skan.AI*

July 2024 – April 2026

Agentic-AI, Robotic Process Automation Discovery Framework

- Modeled a **next activity prediction** model on a custom-built **transformer** framework; fine-tuned open **LLM architectures** (Qwen, Flan-T5, Gemma) to **81%** top-1 accuracy over **200+** activity classes.
- Built rolling **automatability scores** at task sub-sequence level using **perplexity**, surfacing **35+** automatable candidate patterns per client process and letting clients tune frequency factor and pattern length.

Agentic-AI, Automatic Playbook Discovery and Agentic Execution Framework

- Built an **end-to-end process understanding pipeline** for automatic playbook discovery from raw event logs and user interaction data.
- Leveraged **Ollama**-served open-source models in a multi-model **RAG** system for contextual reasoning and action planning, orchestrating **agentic execution workflows** over structured process intelligence.
- Productionized inference through **ONNX Runtime** export with operator fusion and **INT8 quantization** — **3.8×** smaller memory footprint and **41%** lower p95 latency.
- Accelerated decoding via **speculative decoding** (draft-target verification) and **KV-cache** optimizations — paged attention, prefix reuse across turns, continuous batching — for **2.9×** serving throughput and **45%** lower GPU cost at equal SLA.

Data Scientist Intern, *Skan.AI* — Event Data Extraction using LLMs

May 2023 – Jul 2023

- Developed fine-tuned **BERT**, **BART**, **Pegasus** models for event data extraction using **Teacher-Student** distillation, yielding domain-specific variants **6×** smaller at **98%** of teacher F1.
- Applied **clustering algorithms** to improve consistency and accuracy of extracted event data.

Publications

Glaucoma Abnormality Classification using YOLO and SWIN Transformer

May 2024

Madukuri Shaurya, Anumeha Varma, Monika Agrawal — [10.1109/ISBI56570.2024.10635301](https://doi.org/10.1109/ISBI56570.2024.10635301)



Projects

PaCT: Parallel Continuous Thought for Reasoning in LLMs

- Proposed a parallel latent reasoning framework removing **recurrent unrolling** in continuous thought, restructuring reasoning as stage-wise latent resynthesis via a **learnable query codebook + cross-attention**, refined by bidirectional within-stage interaction under causal decoding.
- Achieved **4–9×** **training speedup** over Coconut while matching or improving accuracy on GSM8k, ProntoQA and ProsQA; scaled to **LLaMA-3B/8B** with a consistent 5× reduction.

Forge — Custom GPU Kernels for LLM Training Optimization

- Built custom **Triton GPU kernels** for memory-efficient LLM training, including a **fused linear + cross-entropy** kernel with tiled vocabulary computation that avoids full logit materialization, cutting peak activation memory **62%** at 128k vocabulary.
- Developed fused **RMSNorm/LayerNorm** (single-pass reduction + normalization), **SwiGLU** and **RoPE** kernels with fused forward and backward passes, **1.4×** faster than the PyTorch baseline.

Fully Homomorphic Encryption Scheme

- Implemented **lattice-based short integer solutions**, **Regev's scheme** using learning with errors (LWE), and **fully homomorphic encryption** over ring-LWE.