Pramod Chunduri

pramodc@gatech.edu ◇ LinkedIn ◇ +1-404-980-8546 ◇ https://pchunduri6.github.io

Passionate about building efficient data management systems for AI applications. Current research focuses on building cost-effective and accurate data management systems for applications powered by large language models (LLMs). Track record of accelerating complex video analytics queries, such as action localization, by up to 20× using innovative systems optimizations.

Education

Ph.D. in Computer Science *Aug 2018 - Present*
Georgia Institute of Technology, Advisor: Prof. Joy Arulraj

B.Tech. in Computer Science and Engineering *May 2017*
Indian Institute of Technology, Kanpur

Research Experience

Video Analytics Systems *Georgia Tech* *Spring 2019 - Present*

- *SketchQL*: Developed a novel visual querying system to retrieve complex video events from user sketches. Leveraged transformers-based similarity search to improve retrieval accuracy by 20 F1 points over contrastive learning methods.
- *Zeus*: Built a novel reinforcement learning-driven video analytics system to rapidly localize complex actions like accidents in videos. Improved query throughput by up to 22× over the state-of-the-art.
- *Tracer*: Developed an innovative VDBMS for efficiently re-identifying video objects, using adaptive query processing and a probabilistic search model, surpassing the state-of-the-art system by 3.9× on diverse datasets.

Open-source AI/LLM contributions *Georgia Tech* *2023*

- *rag-demystified*: Built an advanced retrieval-augmented generation (RAG) framework for complex question-answering over documents. Demystified the inner workings of multi-hop RAG/LLM agents using a low-dependency codebase. Gained significant recognition in the tech community, accruing 570 GitHub stars and trending on Hacker News.
- *stargazers-reloaded*: Built an LLM-powered application to analyze GitHub communities. Implemented novel optimizations, including LLM model cascades and Batch Prompting, to reduce API costs by up to 10×.
- *EvaDB*: Actively contributing to EvaDB, a popular database system for AI-powered applications.

Publications

1. SketchQL: Video Moment Querying with a Visual Query Interface, *SIGMOD 2025.*
   Renzhi Wu*, Pramod Chunduri*, Ali Payani, Xu Chu, Joy Arulraj, Kexin Rong.
2. Tracer: Efficient Object Re-Identification in Networked Cameras through Adaptive Query Processing, *Preprint.*
   Pramod Chunduri, Yao Lu, Joy Arulraj.
3. EVA: An End-to-End Exploratory Video Analytics System, *DEEM at SIGMOD 2023.*
   Gaurav Tarlok Kakkar, Jiashen Cao, Pramod Chunduri, Zhuangdi Xu, Suryatej Reddy Vyalla, and 6 others.
4. Seiden: Revisiting Query Processing in Video Database Systems, *VLDB 2023.*
   Jaeho Bang, Gaurav Tarlok Kakkar, Pramod Chunduri, Subrata Mitra, Joy Arulraj.
5. Zeus: Efficiently Localizing Actions in Videos using Reinforcement Learning, *SIGMOD 2022.*
   Pramod Chunduri, Jaeho Bang, Yao Lu, Joy Arulraj.

Technical Skills

Coursework and Expertise	Video Analytics, Systems for AI, Large Language Models (LLMs), Database Systems, Efficient AI, Computer Vision, Data Management, Reinforcement Learning, Computer Systems, Video Understanding, Deep Learning, Intelligent Systems
Programming Languages	C, C++, Python, SQL, bash
Research tools	PyTorch, OpenCV, Scikit-learn, OpenAI Gym, MySQL, Jupyter, LaTeX

Internships

Research Scientist Intern, Adobe Research *San Jose, USA* *Summer 2022*

- Led the development of a multi-stage pipeline for learning user-specific video edits, utilizing a 3D Resnet-based model feature extractor and K-Means algorithm for precise content-based clustering.
- Advanced video editing efficiency by recommending personalized edits to the users with 80% accuracy.

Visiting Research Intern *Advisor: Prof. Boris Grot, The University of Edinburgh* *July 2017 - Jan 2018*

- Empirically established the memory-bound nature of graph workloads through profiling with native performance counters. Integrated software prefetching in graph applications to hide memory latency and accelerate graph analytics.

Flipkart *Security Engineering Team, Bangalore, India* *Summer 2016*

- Implemented automated vulnerability detection and static source code analysis with end-to-end security report.

Awards

- School of Computer Science Incubator Graduate Fellowship, Co-recipient, Georgia Tech, Fall 2021.
- Microsoft Azure Grant of $20,000 research compute credits, Fall 2021 - Spring 2022.
- Chair's fellowship, School of Computer Science, Georgia Tech, 2018.
- All India Rank 216 in JEE Advanced, All India Rank 107 in JEE Mains, All India Rank 3 in VITEEE, 2013.
- Kishore Vaigyanik Protsahan Yojana (KVPY) fellowship, DST, Government of India, 2012.
- State-wise top 1% certificate, National Standard Examination in Physics (NSEP), IAPT, India, 2012.

Teaching

Project Mentor *CS6235 Real-Time Embedded Systems, Georgia Tech* *Spring 2021 - Spring 2023*

- Mentored a team of 20 graduate students over 2 years to build the non-profit research platform NORP. Led the development of 1) Data visualization and analytics platform, 2) Non-profit data integration with social science researchers.

Graduate Teaching Assistant *CS6422 Database System Implementation, Georgia Tech* *Fall 2020, Fall 2021*

- Served as the co-head TA for a 60-student class. Designed core database systems projects (Buffer Manager, B+-Tree, Query Execution) as part of the educational BuzzDB database, Organized lectures, Designed Exams.

Core Team Member *Counselling Service, IIT Kanpur* *Fall 2015 - Fall 2017*

- Supervised Junior Academic Mentors to organize remedial classes, provided one-on-one remedial academic assistance for freshmen and sophomores.

Selected Academic Projects

Efficient 3D Object Detection *Advanced Computer Vision, James Hays, Georgia Tech* *Fall 2020*

- Accelerated 3D object detection on LiDAR Point Clouds by up to 12% using conditional compute blocks. Used a multi-layer perceptron model to enable the conditional usage of shallow neural networks for *simple* scenes. Report.

Automatic Video Labeling *Data Management and Machine Learning, Xu Chu, Georgia Tech* *Fall 2019*

- Designed a video labeling framework to automatically generate training data for action recognition. Utilized an LSTM network to encode similarity between video clips. Used a hierarchical generative model for class inference. Report.

Action recognition using Model Specialization *Deep Learning, Dhruv Batra, Georgia Tech* *Fall 2019*

- Designed a specialized CNN-RNN Encoder-Decoder architecture to accelerate action recognition in videos. Used the specialized architecture to improve action recognition throughput by 80× over baselines. Report.