

Gabriel Paulos

gabrielp7700@gmail.com | gabrielpaulos.com | linkedin: gabriel-paulos | github: gabriel-paulos

Education

University of Toronto, BSc in Computer Engineering Sept 2018 – May 2024

- Dean's Honor List 2022, 2023

Research Experience

Department of Computer Science, University of Toronto Aug 2024 – Present

Research Assistant, Advised by Prof. Xujie Si

- Researching neuro-symbolic proofs for object motions in a scene
- Conducted a review of multimodal vision systems and neuro-symbolic vision systems
- developing a domain-specific language to capture motion trajectories through search-based program synthesis
- The goal of the project is to improve the interpretability of a machine learning system

University of Toronto, Middleware Systems Research Group (MSRG) May 2023 – Dec 2023

Research Assistant, Advised by Prof. Hans-Arno Jacobsen

- Co-designed and co-developed software, in this case middleware, that enables data transfer for workloads in large cloud data processing systems with separate computing and storage nodes
- Built scripts that enabled the creation of jobs to handle large workloads (>100 GB) using Kubernetes and Docker, improving processing efficiency by 50%
- Assisted a graduate student with the evaluation of their research for their thesis that manages memory resources for large workloads

Publications

Efficient Data Transfer in Shared-storage Cloud Data Processing Systems with Optics Oct 2023

Gabriel Paulos, Tongkun Zhang, Yuqiu Zhang, Gerry Zhu, Hans-Arno Jacobsen,
10.5555/3615924.3623630

Work Experience

LyricChief May 2022 – Present

Software Engineer/Founder

- Co-founded a start-up dedicated to creating a vibrant community for sharing the diverse stories and cultural heritage of African music, engaging 1,000s of monthly users and fostering cross-cultural connections
- Spearheading the development of a RESTful API using Ruby on Rails and GraphQL to fetch data to respond to requests, implemented within a CI/CD pipeline via GitLab CI, to empower streaming platforms with seamless access to our lyrics and annotations
- Played a key role in shaping design and business strategies, taking the lead on technology stack selection in a dynamic agile environment, ensuring alignment with project goals and user needs
- Architected and deployed a high-performance PostgreSQL database to efficiently manage and store lyrics, lyric annotations, and song metadata, ensuring scalability and reliability for user access

Projects

Hallucination Mitigation of VLMs using MCTS DPO

- The purpose of the project is to introduce an improvement on the consistency and accuracy of VLMs in answering multi-step and complex VQA tasks
- Problem: Current VLM frameworks do poorly on complex tasks that requires it to perform multi-step reasoning, because of its training objective and the richness of the video modality
- Proposed Solution: Using a two-part system, we use MCTS to collective preference data and break down rewards step wise, then use DPO to update the policy at each newly generated step level and improve the reasoning ability of the VLM

Capstone Project, Automatic Audio/Video Content Editing

Github

- The purpose of the project was to identify and subsequently edit out frames (both audio and visuals) that contained filler words, such as “um” or “you know”, from videos
- Collected and co-created a custom dataset for testing with 50GB of data
- Led design of the pipeline and training of the neural networks (BERT, wav2vec)
- Researched influential papers in the fields of speech recognition technology and Natural Language Processing (NLP)
- Our design was able to achieve a 90% success rate in correctly detecting filler words and editing out the frames that contained them in the project

Reinforcement Learning for Autonomous Vehicles Class Project

Github

- Re-implementation of Dreamer-RNN model
- Applied entropy regularisation on the actor model in order to improve the policy
- Added functionality to the environment, throttle-brake dead zone scheduling and brake speed threshold, to improve training

University of Toronto Hatchery Startup, Depactor

- Led deployment of the logistic regression machine learning model that was the main component of Depactor’s product
- Developed the preprocessed tokenizing function that was fed as input to the model
- Created a program to web scrape Enron’s email corpus and divide it into testing, validation and training sets
- Presented product and its capabilities to group of investors, qualifying as finalists for the Hatchery’s 2020 demo day

Mentorship/Community Service

High School & University Tutor

Jan 2024 – Present

- Prepared & led tutoring sessions by preparing topics on various concepts in AI and CS level for high school students
- Prepared & led academic review sessions for first year engineering students

Technologies

Languages: Python, C/C++, MATLAB, Bash, Java

Tools: Git, Kubernetes, Docker, Linux/Unix, Blender

Libraries: PyTorch, Tensorflow, NumPy, OpenCV, scikit-learn, SciPy, Jupyter, Matplotlib