

Yi Lu

(+1) 437-605-1024 | tomlu@cs.toronto.edu | LinkedIn | Github | Google Scholar

EDUCATION

University of Toronto

Master of Science in Applied Computing (AI Concentration)

Current GPA: 4.0/4.0

Sept.2024 – Dec.2025 (Expected)

University of Manchester

B. Sc. Hons Computer Science (Specialized in Artificial Intelligence)

GPA: 4.0/4.0 (First Class Honours)

Sept.2020 – Jun.2023

PUBLICATIONS

RSVP: Reasoning Segmentation via Visual Prompting and Multi-Modal Chain-of-Thought

ACL 2025 Main

- First-authored and Project Lead, responsible for model implementation, experiment and paper writing.
- Proposed a novel **State-of-the-art** (SOTA) Training-free Reasoning Segmentation Model, unleashed Vision Language Model's localization capability by utilizing region-aware multi-modal visual prompt and hierarchical reasoning architecture.
- Implemented solution surpassed previous SOTA fine-tuned models by up to **9.2 cIoU metric** on ReasonSeg dataset.

VEU-Bench: Towards Comprehensive Understanding of Video Editing

CVPR 2025 Highlight

- Accepted as CVPR 2025 Highlight, responsible for dataset curation, benchmarking and paper writing.
- Proposed a comprehensive fine-grained benchmark dataset for high-level video editing comprehension, demonstrated research gap by evaluating **11 SOTA video-language models** on high-level video editing comprehension tasks.
- **Fine-tuned Qwen2-VL 7B** as baseline model using **ModelScope Swift framework**. Resulting model achieved **GPT-4o-level performance** on video editing understanding tasks.

Video Repurposing from User Generated Content: A Large-scale Dataset and Benchmark

AAAI 2025 Poster

- Paper accepted as **AAAI 2025 poster**, Contributed to the dataset cleaning, construction and paper writing.
- Proposed cross-modality video long-to-short (Video Repurposing) task, curated a large-scale dataset and a baseline model.

RESEARCH INTERNSHIPS

Research Intern

Mar.2025 - Present

TIGER Lab @ University of Waterloo

Topic: Reinforcement Learning, LLM Agent

- Under supervision of Prof. Wenhui Chen, currently working on Reinforcement Learning and agentic LLMs.
- Core contributor to **Verl-Tool** framework, responsible for Python Code Interpreter tool, Agentic Search tool development and integration, Tool-aware LLM Coder model training and Agentic WebSearch LLM training.

Applied Research Intern

May.2025 - Present

ModiFace (L'Oréal's AI Lab)

Topic: Controllable Video Generation.

- Responsible for developing a training-free trajectory control framework for LVDMs, work is in submission to AAAI 2026.
- Responsible for developing a portrait-aware test-time training motion customization method capable of producing photorealistic portrait try-on videos given an instructive prompt and a portrait image. Work is currently in progress.

Applied Research Intern

Jan.2025 - May.2025

Opus AI Research

Topic: Multi-Modal Large Language Models

- Responsible for developing and benchmarking a training-free reasoning segmentation model. In charge of project leading, model implementation, experiments, and paper writing.
- Responsible for constructing a benchmark dataset for high-level video editing comprehension. Repurposed and categorized mainstream VQA datasets for video editing comprehension tasks. Benchmarked 11 SOTA MLLMs on the curated dataset and contributed to paper writing, fine-tuned Qwen2-VL 7B as the baseline model.

COMPETITIONS

Multilingual Video Reasoning Evaluation Challenge @CVPR 2025.

Winner

Complex Video Reasoning and Robustness Evaluation Challenge @CVPR2025.

Runner-up

Long-Term VideoQA challenge of the LOVEU Workshop @CVPR2024.

Winner

Hour-long VideoQA challenge of the Second Perception Test challenge @ECCV2024.

Runner-up

INDUSTRIAL INTERNSHIPS

Machine Learning Engineer

Apr.2024 - Jan.2025

OpusClip (a16z Top50 GenAI Startup)

- Developed **Clip-Copilot**, an LLM-driven interactive video editing agent later evolved into core product **Clip Anything**.
- Designed **MM-Screenplay**, a multi-modal **video understanding** framework leveraging **WhisperX**, **GPT-4** and **Gemini** for retrieving visual and contextual information from **hour-long** videos based on user queries.

Machine Learning Engineer

Dec.2023 - Mar.2024

FaceMind (AI Startup delivering Customized Local-Deployable Chatbots)

- Implemented a **RAG system** using **Milvus** vector database and **LangChain**, enabling **long-term memory**, **context recall**, and **consistent dialogue** for LLM-based chatbots, expanded retrieval chat history length by **up to 200x**.
- **Fine-tuned LLMs** (LLaMA2-7b, Mixtral-8x7b, Qwen-14b, etc.) into Chatbots using **QLoRA** with **LLaMA-Factory**, characterized models' self-awareness and dialogue tone, injected business-specific domain knowledge to the LLMs.

Data Scientist

Aug.2023 - Dec.2023

ByteDance (Parent Company of Tiktok)

- **Fine-tuned an proprietary Visual Language Model** using **LoRA** with **Huggingface-PEFT** for commercial intent understanding, deploying it to identify and analyze business intent in short video content.
- Built a **zero-shot image-matching service** leveraging fine-tuned **CLIP**, **OWL-ViT**, **YOLOv8**, **OpenCV** and **Selenium**, capable of matching products in short videos with images crawled down from related advertising websites.

TECHNICAL SKILLS

Programming Languages: Python, Java, C/C++, SQL

Machine Learning: PyTorch, Huggingface Transformers, PEFT, Llama-Factory, Swift, NumPy, Pandas, Scikit-learn, Tensorflow

Platform, Libraries and Tools: Linux, Git, Conda, Azure, Docker, MySQL, Milvus, Google Cloud Platform