Yuxiang Qiu

Email: yuxqiu@proton.me | GitHub: @yuxqiu | Website: yuxqiu.github.io

Education

University College London 09/2021 - 06/2025

MEng Computer Science

- **Grades**: 1st class (87%, 1st year, rank: 1/150), 1st class (86%, 2nd year)
- **Coursework**: Algorithms for Computer Systems, Computer Architecture & Concurrency, Intelligent Systems, Intro to Cryptography, Logic, Malware, Networked Systems, Security, Supervised Learning, Theory of Computation
- **Thesis**: Work on ZKP for efficient blockchain light client. Advised by Prof. Philipp Jovanovic and Alberto Sonnino.

Georgia Institute of Technology 08/2023 – 05/2024

BS Computer Science (Exchange Student)

- **GPA**: 4.0/4.0
- **Coursework**: Blockchain & Cryptocurrency, Compiler & Interpreter, Computer Graphics, Deep Learning, Design & Analysis of Algorithm, Processor Design, Quantum Computing, Zero Knowledge Proofs (S2023 MOOC, self-taught)

Experience

Thesis: Trustless Efficient Light Clients Made Practical 10/2024 - Present

- Background: Light clients are an important part of the blockchain ecosystem. Many light client protocols currently exist in different blockchains. However, they are 1) resource-intensive, as the data to be downloaded and the operations to be performed are sublinear or linearly related to the chain size; 2) not generalized and tied to specific blockchains; and 3) inefficient for the provers. We propose to use folding-based SNARK to solve these issues.

Research Assistant 06/2024 – 09/2024

UCL Software Optimisation, Learning and Analytics Research Lab *London, UK*

- Background: Recent advances in LLM show the promise of using it to judge text quality. However, current methods lack interpretability and are vulnerable to adversarial attacks. To solve these, we propose *TaskEval*, a method to score an explanation by measuring how well an LLM can accomplish tasks with this explanation.
- Research: reviewed 10+ available datasets, proposed LLM-as-a-judge as the **baseline**, evaluated and **enhanced 4 text perturbation methods**, designed ways to **improve and measure the diversity** of generated text
- Implementation: **integrated SWE-bench** into the eval framework, implemented fault localization and differential testing evaluator tasks, designed the dynamic transitivity-based comparison algorithm used in surveys
- Experiment: designed and conducted experiments to **analyze the performance** (in terms of agreement, Kendall's Tau, and Spearman's correlation) of TaskEval in different settings (with CoT, different perturbations, etc.)
- Advisor: Prof. Federica Sarro and Prof. Sergey Mechtaev

Software Development Engineer Intern 06/2023 – 08/2023

Amazon *London, UK*

- Researched cross-platform portability of Java apps running on Windows, resulting in a ~10-page research report
- Delved into the Java SE Specifications (JVMS and JLS), the JAR file specifications, and the OpenJDK source code
- Developed a Java application and library that performs **incompatibility detection at the bytecode** level (checking for 7 different types of cross-platform issues) with **~80% accuracy and 90%+ recall**
- Optimized libraries by profiling hot spots and bringing parallelism to CPU-bound tasks, resulting in a **3x speedup**

Teaching Assistant

UCL

- 2024-2025: COMP0002 Principles of Programming, COMP0004 Object-Oriented Programming
- 2022-2023 Programming Tutor: Tutored 12 students in 6 programming languages (C, C++, Rust, Haskell, Java, Python) and familiarized them with shell scripting, computer networking, and frontend/backend development

Open Source Contributions

- **AI**: `pytorch/torcheval` (#195), `princeton-nlp/SWE-bench` (#186, #189, #212)
- **PL**: `rust-lang/rust-clippy` (#11865, #12084, #12094), `typst/biblatex` (#34)

Projects

TrueLearn 01/2023 - 08/2023

- Led a team of 4 students to **implement a Python machine-learning library** with a family of baseline and Bayesian classifiers for building learner models to predict their engagement with educational resources
- **Created 9 static and interactive visualizations** to present the learner representations in humanly-intuitive ways
- Conducted **hyperparameter tuning** via grid search and evaluated library scalability by analyzing wall-clock time
- Augmented the PEEKC dataset (with 30000+ Wikipedia data) to provide richer info during the entity linking process
- Advisor: Dr. Sahan Bulathwela

Logic Parser 10/2022 - 12/2022

- Devised a **one-pass iterative parser** and a **tableau-based SAT solver** for propositional and predicate logics
- Built efficient iterative algorithms for AST operations that support processing logic formulas of arbitrary size in a scalable way, with **performance comparable to the SOTA z3 solver** for propositional logic

Awards

UCL Studentship for Research 2024
UCL Faculty Undergraduate Scholarships for Excellence (1 student per faculty, 1 out of 1000+ students) 2022

Publications

TaskEval - Using LLMs to Evaluate Natural Text Artifacts: A Case Study on Patch Explanations
David Williams, Yuxiang Qiu, Peichu Xie, Sergey Mechtaev, Federica Sarro, Mark Harman
In Preparation

A Toolbox for Modelling Engagement with Educational Videos
Yuxiang Qiu, Karim Djemili, Denis Elezi, Aaneel Shalman Srazali, Mar'ia P'erez-Ortiz, Emine Yilmaz, John Shawe-Taylor and Sahan Bulathwela
Proceedings of the AAAI Conference on Artificial Intelligence, 2024

TrueLearn: A Python Library for Personalised Informational Recommendations with (Implicit) Feedback
Yuxiang Qiu, Karim Djemili, Denis Elezi, Aaneel Shalman, María Pérez Ortiz and Sahan Bulathwela
6th Workshop on Online Recommender Systems and User Modeling, ACM RecSys 2023

Skills

Languages: C++, C, Python, Rust, Java, Verilog, Solidity, HTML, CSS, JavaScript, Haskell, x86 Assembly, GLSL
Libraries: ANTLR, arkworks, Bootstrap, Flask, Koa, openai, OpenCV, OpenGL, PyTorch, scikit-learn, Vue.js