Fabrice Harel-Canada

Researcher · Software Engineer

(+1) 954-805-5735 | fabricehc@cs.ucla.edu | fabrice.harel-canada.com | fabriceyhc

Biography

Recent UCLA Computer Science PhD specializing in pioneering robust and holistic evaluation frameworks for AI systems. My doctoral research includes award-winning work on human-centric LLM quality assessment, novel methods in AI testing, data augmentation (e.g., sibylvariant transformations), content provenance (watermarking), and most recently NLP applied to substance use epidemiology.

Research Experience

The Wellness Center @ LA General Hospital

Supervisor: Rebecca Trotzky

Applied Research Engineer — *July 2025 - Present*

- Engineer and deploy fine-tuned LLMs to extract substance use history from clinical notes, achieving high-precision classification between active and historical usage.
- Architected a large-scale data pipeline to perform epidemiological statistical analysis on 1M+ electronic health records to identify longitudinal trends in substance use.

University of California, Los Angeles

Supervisor: Chelsea Shover

Graduate Student Researcher — *Dec 2023 - July 2025*

- Developed a context-aware drug misuse detection system (F1 > 0.95), significantly reducing false positives in clinical screening across multiple drug classes.
- Helped engineer a predictive geospatial model to forecast overdose hotspots in Los Angeles County, integrating multi-modal data to inform public health intervention strategies.

University of California, Los Angeles

Supervisor: Miryung Kim

Graduate Student Researcher — *June 2019 - Dec 2023*

- Validated the real-world robustness of AI watermarking through large-scale experimentation, overturning theoretical assumptions about watermark erasability and demonstrating a 90% retention rate against state-of-the-art attacks.
- Developed a theoretically grounded framework to assess the psychological depth of LLMs, revealing that modern LLMs match or exceed the abilities of advanced human short story writers.
- Created automated evaluation metrics for natural language generation systems, including measures of linguistic diversity – e.g. semantic, syntactic, morphological, phonological – and as well as general engagement.
- Designed and open-sourced a new class of data augmentations called sibylvariant transforms for enhancing NLP model robustness.
- Assessed the reliability of structural testing metrics, such as neuron coverage, in evaluating deep learning models.
- Developed diversity-enhancing extensions for adversarial attack algorithms to improve robustness testing.

University of California, Los Angeles

Supervisor: Judson Caskey

Research Assistant — *July 2020 - Oct 2020*

- Leveraged BERT-based transformers and LDA to execute sentiment analysis and topic modeling on a massive dataset of 1M+ Amazon and Glassdoor reviews.

Professional Experience

Kairno

Founder and CEO — *May 2025 - Jan 2026*

- Developed and deployed a full-stack system (@ `kairno.com`) that automates rigorous academic peer review, ensuring compliance with field-specific reporting guidelines and best practices.

Verox Tech

Supervisor: Michael Simhai

Project Manager + Developer Team Lead — *July 2015 – Sept 2018*

- Spearheaded the end-to-end development and launch of a cloud-native insurance platform, managing a cross-functional team of 23 engineers and designers.
- Architected scalable C# Web APIs and optimized complex T-SQL procedures, maintaining high performance for high-concurrency database transactions.
- Led a mission-critical migration of legacy data across three disparate systems, ensuring 100% data integrity and zero downtime during the transition.
- Responsible for producing, coordinating and reviewing nearly every aspect of the SDLC: code reviews and code optimization, requirements and spec generation, database design, technical documentation, UI design standards, and deployment scripts.

People's Trust Insurance *Supervisor: Scott Thompson*

Business + Data Analyst *Dec 2012 – Sept 2018*

- Served as lead data strategist for the executive suite, delivering 100+ high-impact analytical reports that drove multi-million dollar business decisions.
- Developed numerous front-end / back-end application features that were successfully deployed into production to automate tedious and time-intensive workflows.
- Bridged the gap between business stakeholders and engineering teams across 60+ agile sprints, overseeing the transition to a modernized policy and claim management ecosystem.

Cooperative Living Organization

President *Dec 2010 – May 2012*

- Directed organizational operations by managing 17 department heads and implementing standardized policy frameworks.
- Optimized revenue streams and marketing, increasing house occupancy from 75% to 100% capacity and driving a 57% increase in net income.

Publications

[9] **Fabrice Harel-Canada***, Boran Erol*, Connor Choi, Jason Liu, Gary Jiarui Song, Nanyun Peng, and Amit Sahai. "Sandcastles in the Storm: Revisiting the (Im)possibility of Strong Watermarking". In: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar. Vienna, Austria: Association for Computational Linguistics, July 2025, pp. 29698–29735. ISBN: 979-8-89176-251-0. DOI: 10.18653/v1/2025.acl-long.1436. URL: https://aclanthology.org/2025.acl-long.1436/.

[8] **Fabrice Harel-Canada**, Anabel Salimian, Brandon Moghanian, Sarah Clingan, Allan Nguyen, Tucker Avra, Michelle Poimboeuf, Ruby Romero, Arthur Funnell, Panayiotis Petousis, Michael Shin, Nanyun Peng, Chelsea L. Shover, and David Goodman-Meza. "Enhancing substance use detection in clinical notes with large language models". In: *Drug and Alcohol Dependence* 276 (2025), p. 112888. ISSN: 0376-8716. DOI: https://doi.org/10.1016/j.drugalcdep.2025.112888. URL: https://www.sciencedirect.com/science/article/pii/S0376871625003412.

[7] Arthur J. Funnell, Panayiotis Petousis, **Fabrice Harel-Canada**, Ruby Romero, Alex A. T. Bui, Adam Koncsol, Hritika Chaturvedi, Chelsea Shover, and David Goodman-Meza. *Improving Drug Identification in Overdose Death Surveillance using Large Language Models*. 2025. arXiv: 2507.12679 [cs.CL]. URL: https://arxiv.org/abs/2507.12679.

[6] Rohan Wadhawan, **Fabrice Harel-Canada**, Zi-Yi Dou, Suhaila Shakiah, Robinson Piramuthu, and Nanyun Peng. "VaPR–Vision-language Preference alignment for Reasoning". In: *COLM 2025 Conference*. 2025. URL: https://arxiv.org/abs/2510.01700.

[5] **Fabrice Harel-Canada**, Hanyu Zhou, Sreya Muppalla, Zeynep Senahan Yildiz, Miryung Kim, Amit Sahai, and Nanyun Peng. "Measuring Psychological Depth in Language Models". In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Ed. by Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen. **Outstanding Paper Award (Top 0.25% of submissions)**. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 17162–17196. DOI: 10.18653/v1/2024.emnlp-main.953. URL: https://aclanthology.org/2024.emnlp-main.953/.

[4] Hong Jin Kang*, **Fabrice Harel-Canada***, Muhammad Ali Gulzar, Nanyun Peng, and Miryung Kim. "Human-in-the-Loop Synthetic Text Data Inspection with Provenance Tracking". In: *Findings of the Association for Computational Linguistics: NAACL 2024*. Ed. by Kevin Duh, Helena Gomez, and Steven Bethard. Mexico City, Mexico: Association for Computational Linguistics, June 2024, pp. 3118–3129. DOI: 10.18653/v1/2024.findings-naacl.197. URL: https://aclanthology.org/2024.findings-naacl.197/.

[3] **Fabrice Harel-Canada**, Muhammad Ali Gulzar, Nanyun Peng, and Miryung Kim. "Sibylvariant Transformations for Robust Text Classification". In: *Findings of the Association for Computational Linguistics: ACL 2022*. Ed. by Smaranda Muresan, Preslav Nakov, and Aline Villavicencio. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 1771–1788. DOI: 10.18653/v1/2022.findings-acl.140. URL: https://aclanthology.org/2022.findings-acl.140/.

[2] Guangxuan Xu, Ruibo Liu, **Fabrice Harel-Canada**, Nischal Reddy Chandra, and Nanyun Peng. “EnDex: Evaluation of Dialogue Engagingness at Scale”. In: *Findings of the Association for Computational Linguistics: EMNLP 2022*. Ed. by Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, Dec. 2022, pp. 4884–4893. DOI: 10.18653/v1/2022.findings-emnlp.359. URL: https://aclanthology.org/2022.findings-emnlp.359/.

[1] **Fabrice Harel-Canada**, Lingxiao Wang, Muhammad Ali Gulzar, Quanquan Gu, and Miryung Kim. “Is neuron coverage a meaningful measure for testing deep neural networks?” In: *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 2020, pp. 851–862.

A more comprehensive list is available at: Semantic Scholar Profile

Education

University of California, Los Angeles (UCLA) *Advisor: Miryung Kim*

PhD | Computer Science *Jan 2020 - June 2025*

- 3.92 / 4.00 GPA
- Majors: Artificial Intelligence + Software Engineering | Minor: Data Science Computing
- Thesis: *From Metrics to Meaning: Advancing Evaluation Frameworks for Robust and Human-Centric AI*

University of California, Los Angeles (UCLA) *Advisor: Miryung Kim*

MS | Computer Science *Sept 2018 - Dec 2019*

- Major: Artificial Intelligence
- Thesis: *Is Neuron Coverage a Meaningful Measure for Testing Deep Neural Networks?*

University of Florida (UF)

BS | Business Administration *Aug 2008 - May 2012*

- Majors: Management + Information Systems and Operations Management
- Minors: Entrepreneurship + Finance
- Thesis: *The Viability of the Student-Run Housing Cooperative: A Case Study of the Cooperative Living Organization*

Skills

Programming	python, sql, java, c#, linq, html, css, js, scala, problog, prolog, LaTeX
Frameworks	huggingface, pytorch, tensorflow, keras, spark, numpy, pandas, scipy, sklearn, opencv, matplotlib, seaborn
Tools	ssms, ssis, ssrs, tfs, azure devops, visual studio, excel, aws, gcp, exchange server, moqups