KIELEH NGONG IVOLINE CLARISSE

Research Interest: Privacy-preserving ML, Provable Fairness, Deep Learning
Email: ivolinengong@gmail.com , **Phone:** +1 802-324-1797
Github, LinkedIn, Website

EDUCATION

PhD Student in Computer Science, University of Vermont, USA *Aug 2021 - Present*
Data Privacy and Provable Fairness
Masters in Computer Engineering, Konya Technical University, Turkey *Sep 2018 - July 2021*
Specializing in Machine Learning and Deep Learning
Bachelors in Computer Engineering, University Of Buea, Cameroon *Sep 2012 - Aug 2016*
Specializing in Software Engineering

PUBLICATIONS

Peer Reviewed

- Protecting Users From Themselves: Safeguarding Contextual Privacy in Interactions with Conversational Agents (**Ngong, I.**,Kadhe, S.,Wang, H.,Murugesan, K., Weisz, J., Dhurandhar, A., Ramamurthy, K.) *2025*
  Association for Computational Linguistics (ACL - Findings)
- SoK: Usability Studies in Differential Privacy (Dibia, O., Stenger, B., Baldasty, S., Bates, M.,**Ngong, I.**, Feng, Y., Near, J.) *2025*
  Privacy Enhancing Technologies Symposium (PETS)
- Differentially Private Learning Needs Better Model Initialization and Self-Distillation (**Ngong, I.**, Near, J., Mireshghallah, N.) *2025*
  Published in Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL)
- OLYMPIA: A Simulation Framework for Evaluating the Concrete Scalability of Secure Aggregation Protocols (**Ngong, I.**, Gibson N., Near, J.) *2024*
  Published in 2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)
- Evaluating the Usability of Differential Privacy Tools with Data Scientists (**Ngong, I.**, K., Stenger, B., Near, J., Feng, Y.) *2024*
  Published in Usable Privacy and Security (SOUPS 2024) colocated with USENIX
- Different Deep Learning Based Classification Models for COVID-19 CT-Scans and Lesion Segmentation Through the cGAN-UNet Hybrid Method (**Ngong, I.**, Baykan, N.) *2023*
  Published in Traitement du Signal Journal
- Feature Extraction Methods for Predicting the Prevalence of Heart Disease (**Ngong, I.**, Baykan, N.) *2022*
  Published in Springer Books Series and presented at **SCA2021** (The Sixth Smart City Applications International Conference)

Peer-Reviewed Workshop Papers

- Protecting Users From Themselves: Safeguarding Contextual Privacy in Interactions with Conversational Agents (**Ngong, I.**,Kadhe, S.,Wang, H.,Murugesan, K., Weisz, J., Dhurandhar, A., Ramamurthy, K.) *2024*
  Socially Responsible Language Modelling Research - **SoLaR at NeurIPS 2024**
- Decreasing Hallucinations in Differentially Private Large Language Models(LLMs) through Self-Distillation (**Ngong, I.**, Near, J., Mireshghallah, N.) *2024*
  Socially Responsible Language Modelling Research - **SoLaR at NeurIPS 2024**
- Evaluating the Usability of Differential Privacy Tools with Data Scientists (**Ngong, I.**, K., Stenger, B., Near, J., Feng, Y.) *2023*
  Theory and Practice of Differential Privacy - **TPDP 2023**

- Distributed HDMM: Optimal Accuracy without a Trusted Curator (Sedimo, R., **Ngong, I.**, K., Near) *2023*
  Theory and Practice of Differential Privacy - **TPDP 2023**
- Towards Auditability for Fairness in Deep Learning (**Ngong, I.**, Maughan, K., Near, J.) *2020*
  Algorithmic Fairness through the lens of Causality and Interpretability Workshop - **AFCI at NeurIPS**

Pre-prints and Others

- How To Audit An AI Model Owned by Someone else (Part 1) (**OpenMined Team)** *2023*
- Prediction Sensitivity: Continual Audit of Counterfactual Fairness in Deployed Classifiers (Maughan,K., **Ngong, I.**,Near, J.) *2022*
- Continual Audit of Individual and Group Fairness in Deployed Classifiers via Prediction Sensitivity(Maughan,K., **Ngong, I.**,Near, J.) *2022*
  Article featured in Montreal AI Ethics Institute
- Maintaining Privacy in Medical Data With Differential Privacy (**Ngong, I.**) *2020*
  Article featured in OpenMined Blog

PUBLICATIONS (IN PROGRESS)

- Towards Privacy in Agentic Workflows (**Ngong, I.**,Kadhe, S.,Murugesan, K., Weisz, J., Dhurandhar, A., Ramamurthy, K.) *Ongoing*
- Differentially Private Multimodal In-context learning (**Ngong, I.**, Near, J.) *Ongoing*
- Differential Private Retrieval-Augmented Generation (RAG) with k-nearest neighbor(KNN) Large Language Models (**Ngong, I.**, Mireshghallah, Near, J., Kamath, G.) *Ongoing*
- MultiUnlearnBench: A Multimodal Benchmark for Selective Information Unlearning in Large Models (Dai,S., **Ngong, I.**, Wu,T.) *Ongoing*
- Compiler Optimization for Reinforcement Learning (**Ngong, I.**, Zhang, W.) *Ongoing*
- Immediate Sensitivity for Differential Privacy (Stevens, T., **Ngong, I.**, K., Near, J.) *Ongoing*
- Backpropagation Clipping for Deep Learning with Differential Privacy *Ongoing*

WORK EXPERIENCE

Research Assistant Vermont, USA
PLAID Lab, University of Vermont *Aug 2021 - Present*

- Advised by Joe Near
- Research on Federated Learning simulation and new differential privacy methods.
- Research on Provable Fairness and Privacy Using Machine Learning, Funded via Amazon Research Award

Research Scientist Intern New York
IBM Research *May 2025 - August 2025*

- Opensourced contextual privacy toolkit - part of a broader effort to support privacy-aware agentic workflows — where decisions about what information to share are guided by the context and purpose of a task, not just static rules about sensitive data.
- Implementing contextual privacy framework (from the last summer) in Granite Guardian - IBM's flagship product for LLM safety.
- Worked on implementing privacy in agentic workflows reporting to Karthikeyan Natesan.
- Studied agentic workflows, Identified research gap, conducted literature review, and developed comprehensive draft proposal on the privacy risks in agentic workflows.
- Creating a benchmark dataset for privacy risks in agentic workflwos which includes synthetic data generation, prompt engineering etc
- Proposed novel metrics for measuring oversharing using a privacy graph.

· Proposed a mitigation approach grounded in contextual privacy, daa minimization and information flow control.

Research Scientist Intern New York
IBM Research *May 2024 - August 2024*

· Worked on Trustworthy Foundation Models (specifically privacy in large language models) reporting to Karthikeyan Natesan.
· Identified research gap, conducted literature review, and developed comprehensive draft proposal on integrating contextual privacy in LLMs.
· Designed and executed experiments including synthetic data generation, prompt engineering, performed user studies etc
· Collected, evaluated analyzed data, tested hypotheses, and refined research direction.
· Collaborating with a team
· Paper from internship accepted at NEURIPS workshop and published in ACL2025
· Got a search-1 patent filing.

Research Scientist Remote
OpenMined *September 2022 - Present*

· Developing a privacy-preserving auditing protocol for AI systems using PySyft "Domain" servers, facilitating precise external audits without exposing additional information beyond what's necessary.
· Finding and implementing new approaches to simulate and evaluate the performance individual privacy accounting in PySyft as part of the DP Team.
· Part of the team organizing the Medical Federated Learning program.
· Co-organizing the OpenMined Research Team.
· Organized and led a study group on privacy-preserving methods.

Independent Researcher Remote
Microsoft Research *June 2022 - Aug 2022*

· Tackle the phase-ordering problem, performing code size and run time reduction in the LLVM compiler using Reinforcement Learning.

Computer Science Research Mentorship Program Fellow Remote, USA
Google *Sep 2021 - Dec 2021*

· Selected as a Google CSRMP Fellow from a highly competitive pool of applicants all over North America to partake in machine intelligence research-field-related discussions in an intimate pod setting with fellow students interested in research while being mentored by a Google researcher.

Learning Technologist California (Remote), USA
DeepLearning.AI *April 2021-July 2021*

· Worked predominantly on the Machine Learning Specialization, Practical Data Science Specialization and Natural Language Processing (NLP) Specialization. Some tasks included:
· Monitoring feedback on live, in-classroom courses; Collecting and organizing issues for improvement.
· Collecting and organizing issues for improvement.
· Git Repo management - creating, fixing, merging, closing issues.
· Creating and fixing issues with quizzes.
· Create supplementary lecture " notebooks" and/or slides to fill in gaps in the lectures.
· Communicate with mentors to identify priorities for improvements.
· Get familiar with content and provide feedback.

QA Specialist California (Remote), USA
DeepLearning.AI *Nov 2020 - Mar 2021*

- Working with the QA and Curriculum Development teams to ensure the quality of course content that is under development as well as the quality of the learner experience in live courses.
- This includes reviewing content and providing feedback as well as managing a pool of volunteer alpha testers (reviewers) during the course development process.
- I collect, organize and prioritize feedback and provide it to the curriculum development team.
- Track various metrics and follow learner feedback to assess content performance and identify areas for improvement of the overall quality of the courses.
- Involved in recruiting, on-boarding and supporting volunteer mentors who support learners in live courses.

Research Intern Vermont, USA
PLAID Lab, University of Vermont *Jun 2020 - Nov 2020*

- Implementing immediate sensitivity for differentially private deep learning in PyTorch.
- Implemented PyTorch-based library for prediction sensitivity.
- Designed and implemented a novel smoothing approach for prediction sensitivity.
- Designed and implemented new experiments on three datasets to evaluate prediction sensitivity as a measure of individual fairness.
- As part of these experiments, implemented an existing fairness mitigation approach and several group fairness metrics to compare against.
- Got final experimental results and wrote the evaluation section of the AFCI paper.
- Wrote & contributed to large parts of the rest of the AFCI paper.
- Designed improvements to the approach & evaluation methodology based on the reviews (this is the ongoing part).

Independent Consultant California (Remote), USA
SkoolMentor *Jun 2020 - Sep 2021*

- Designed and developed curriculum for Machine Learning and Deep Learning for high school students.
- Guided students through design, implementation and publication of research projects in cancer prediction.

Independent Consultant California (Remote), USA
Udacity *Feb 2019 - Feb 2021*

- As a consultant to an educational platform (Udacity), I utilize my specialized knowledge in the fields of AI in HealthCare, Computer Vision, Android Development, Cybersecurity, and Java Development and my strong communication skills to provide mentorship,project reviews and other student support services.

Writer Remote, USA
Towards Data Science (TDS) *Jul 2019*

- Share my thoughts on models that understand and explain human behaviour by writing articles

Junior ML Engineer Remote
Omdena *Nov 2019 - Jan 2020*

- Building AI for Natural Disasters-Optimizing the AfterMath Management of Earthquakes With AI Part of a team that developed a road segmentation model to find the safest route from one point to another in the aftermath of an Earthquake.
- Acquired Geographical Information Systems data for Istanbul from open street maps, created road segmentation masks and built segmentation models using Keras and Tensorflow.
- Covered the majority of the Machine Learning Workflow; from problem definition, data acquisition, data preprocessing, model development and training.

Scholar Remote
Facebook AI Secure and Private Artificial Intelligence Program *May 2018 - Jan 2020*

- Learned about state of the art privacy preserving methods: Federated Learning, Differential Privacy and Encrypted Deep Learning.
- Led a team that organized virtual weekend hackathons on Kaggle. Organized 5 hackathons that tested scholars knowledge on Classification, Sentiment Analysis, Natural Language Processing and Time Series Forecasting - https://github.com/ivyclare/Virtual-Hackathon
- Recognized by the community as an outstanding scholar for creativity, inspiring presence and exceptional work in organizing hackathons.
- Officially recognised as a top 5% (300/6000) scholar, and awarded a Deep Learning Nanodegree sponsored by Facebook AI.

Open Source Contributor Remote
OpenStack *Sep 2018 - Feb 2019*

- Performed search optimization on Storyboard project and documentation changes on the Sahara Project in Python.
- Performed Code reviews and committed patches with gerrit.
- Covered the majority of the Machine Learning Workflow; from problem definition, data acquisition, data preprocessing, model development and training.

Software Developer Buea, Cameroon
Njorku Inc *Feb 2014 - Nov 2016*

- Improved the Njorku search engine and extended the Njorku website
- Built web applications in Java and Php using PlayFramework and Kohanna
- Built android and windows mobile applications.

Software Developer Intern Buea, Cameroom
MegaSoft Sarl *Sep 2015 - Feb 2016*

- Built, designed and tested web and enterprise applications in Java.

OTHER TRAINING

- Participant in the Differential Privacy Summer School, Boston. *2022*
- Participant in the Oxford Machine Learning Summer School *2022*

TALKS

- Breakout session lead in workshop on "DP Beyond Algorithms", at the OpenDP Community Meeting *2024*
- Talk on "Privacy-Preserving Machine Learning", Factored *2023*
- Presented summer research "Compiler Optimizations using Reinforcement Learning", Microsoft Research. *2022*
- "Fairness and Privacy in AI presentation" at the University of Buea *2022*
- Presented our paper "Feature Extraction Methods for Predicting the Prevalence of Heart Disease" at the Sixth Smart City Applications International Conference. *2022*
- Talk at Computer Science Research Day on "Continual Audit of Individual and Group Fairness in Deployed Classifiers via Prediction Sensitivity" at UVM *2021*
- Invited Guest Speaker for **OpenMined Courses** on "Fairness and Privacy Implications of Machine Learning in the Real World" *2021*

· Presented our paper Towards Auditability for Fairness in Deep Learning (**Ngong, I.**, Maughan, K., Near, J.) at the - **AFCI Workshop at NeurIPS** *2020*

· Invited Panelist at the PhD Recruitment Event at UVM *2021*

SERVICE

Program Committee member for the ACM FAccT 2025 *2025*
Ethics Reviewer for NeurIPS2024 *2024*
Program Committee member for the Theory and Practice of DP workshop (TPDP) *2024*
Program Committee member for the Fifth AAAI Workshop on Privacy-Preserving Artificial Intelligence (PPAI-24) *2024*
Ethics Reviewer for NeurIPS2023 *2023*
Program Committee member for Generative AI + Law Workshop at ICML2023 *2023*
Volunteer at the Women in Machine Learning Workshop held at NeurIPS2022 *2022*
Started Differential Privacy reading group at UVM *2022*
Volunteer at the Women in Machine Learning 2020 Workshop held at ICML2020 *2020*
Organizer at Google Developer Group, Buea *2015*
Co-founder, Women Tech Makers, Buea *2015*
Co-founder, Working to Advance African Women (WAAW) Foundation, Buea Stem Cell *2015*

AWARDS/GRANT PROPOSALS

Compiler Optimization using RL Microsoft Reinforcement Learning Open Source Festival Proposal (Awarded $10,000) *2022*

2nd place for best presentation at the CS Research Day *September 2021*
Udacity Scholar *April 2020*
Computer Vision Nanodegree
Facebook Secure and Private AI Scholar *Feb 2020*
Deep Learning Nanodegree
Bertelsmann Data Science Scholar *Jan 2019*
Data Foundations Nanodegree
Google Android Developer Nanodegree Scholar *Mar 2018*
Android Developer Nanodegree Certificate
University of Michigan on Coursera *Oct 2018*
Python Specialization (5 courses)
University of California, San Diego on Coursera *Sep 2016*
Java Specialization (3 courses)

SKILLS

Python	Machine Learning	Data Analysis / Visualization	(Computer Vision)
Java/Android	Deep Learning	Research	Pytorch
Tensorflow	Keras	Matlab	OpenCV

Other Highlights: Scikit-Learn, Numpy, Matplotlib, Pandas, Docker, Github, Gerrit, MySQL, SQLite,Postgresql, Windows, Linux, Leadership & Mentoring, Communication, GANs, NLP