Skander Moalla

skan.moalla@gmail.com | skandermoalla.com | Lausanne, Switzerland
skander.moalla@epfl.ch | skandermoalla | +41 77 259 52 15

EDUCATION

EPFL **Lausanne, Switzerland**
PhD in Computer and Communication Sciences *Sep 2022 - Present - Fall 2026 (expected graduation)*
Deep reinforcement learning (RL) and LLM post-training. Advised by Prof. Caglar Gulcehre (CLAIRE).

University of Oxford **Oxford, UK**
M.Sc. in Advanced Computer Science *Oct 2021 - Sep 2022*
M.Sc. thesis in multi-agent reinforcement learning with Prof. Shimon Whiteson's group (WhiRL).

École Polytechnique **Paris, France**
B.Sc. Mathematics and Computer Science double major *Sep 2017 - Aug 2020*
Graduated summa cum laude | CGPA: 4.09/4.0 | Awarded an excellence scholarship.

PUBLICATIONS

Simon Matrenok*, **Skander Moalla***, and Caglar Gulcehre. Quantile Reward Policy Optimization: Alignment with Pointwise Regression and Exact Partition Functions. ***NeurIPS***, **2025**.

Project Apertus, [...], **Skander Moalla***, [...], Antoine Bosselut, Martin Jaggi, and Imanol Schlag. Apertus: Democratizing open and compliant llms for global language environments. 2025.

Skander Moalla, Andrea Miele, Razvan Pascanu, and Caglar Gulcehre. No Representation, No Trust: Connecting Representation, Collapse, and Trust Issues in PPO. ***NeurIPS***, **2024**.

Xiuying Wei, **Skander Moalla**, Razvan Pascanu, and Caglar Gulcehre. Building on Efficient Foundations: Effectively Training LLMs with Structured Feedforward Layers. ***NeurIPS***, **2024**.

Skander Moalla*, Manuel Madeira*, Lorenzo Riccio*, and Joonhyung Lee*. [Re] Reproducibility Study of Behavior Transformers. ***ReScience C***, **9(2)**, **2023. Outstanding Paper Awards Honorable Mentions**.

Benjamin Ellis, Jonathan Cook, **Skander Moalla**, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob Foerster, and Shimon Whiteson. SMACv2: An Improved Benchmark for Cooperative Multi-Agent Reinforcement Learning. ***NeurIPS***, **2023**.

EXPERIENCE

Google DeepMind **Paris, France**
PhD Student Researcher *Sep 2025 - Dec 2025*
- Hosted by Alexandre Ramé in the Gemma post-training team (Olivier Bachem, Lucas Dixon).
- Worked on leveraging diversity to improve test-time scaling (pass@k) and sampling in reinforcement learning.
- Contributed to the Gemma open-source JAX codebase (Gemma 3n layer and sampler KV cache and utilities).

Quincus **Singapore (Remote)**
Applied Research Intern *Mar 2021 - Aug 2021*
- Deployed a solution leveraging deep reinforcement learning (RL) to optimize middle-mile logistics.
- Built simulators to model multi-mile parcel shipping and trained deep RL prototypes (PPO, DQN).
- Stack: Python (OpenAI Gym & Baselines, LightGBM, pytest), Docker, NSCC High-Performance Computing.

Amazon **Luxembourg**
Software Development Intern *Sep 2020 - Feb 2021*
- Completed full-stack and DevOps sprint tasks within a two-pizza team owning multiple services end-to-end.
- Designed and built a prototype to detect similar and duplicate business metrics supporting Amazon's transportation network worldwide to ease metric discovery and reduce computation costs, leveraging deep learning and NLP techniques with AWS SageMaker.
- Stack: AWS, Python, Java (Dagger, AutoValue, JUnit, Mockito), TypeScript (ImmutableJS, Mocha, Chai).

EPFL - Laboratory for Topology and Neuroscience & Blue Brain Project **Geneva, Switzerland**
Research Intern *Jan 2020 - Apr 2020*
- Bachelor thesis supervised by Prof. Kathryn Hess Bellwald. (Link to thesis).
- Build a framework for neuron simulation and leveraged tools from information theory and algebraic topology (simplicial homology) to formulate new conclusions about the dynamics of neurons forming directed cliques.
- Stack: Python 3 (DIT, JupyterLab, NumPy, SciPy, Brian2), C++.

Ooredoo **Tunis, Tunisia**
Full-Stack Developer Intern *Jul 2018 - Aug 2018*
- Implemented a POC for a custom notification system in a hybrid mobile app featuring profiling functionalities and allowing Ooredoo to adjust the activity of their notification server. Stack: Ionic 3, Node.js, socket.io.

TEACHING

EPFL - LauzHack **Lausanne, Switzerland**
LauzHack Deep Learning Bootcamp Instructor (Deep Reinforcement Learning) *Summer 2024*

EPFL **Lausanne, Switzerland**
Teaching Assistant & Supervisor for MSc Projects and Theses *Sep 2022 - Present*
- Supervised 9 MSc projects and theses.
- Probability and Statistics (MATH-232 Spring 2023), Machine Learning (CS-433 - Fall 2023 - exam preparation), Artificial Neural Networks/Reinforcement Learning (CS-456 - Spring 2024 - project preparation).

ATSM - Tunisian Association of Mathematics **Tunisia**
Mathematics Tutor *Jul 2018 - Present*
- Designed material for Olympiad problem-solving and trained national mathematics team members for the Mediterranean Math Olympiad (MYMC 2018) and the International Mathematical Olympiad (IMO 2020).

SOFTWARE

Core stack and responsibilities

Co-lead the Swiss AI post-training efforts *July 2024 - Present*
- Developed infrastructure and offline post-training code to train the Swiss open models Apertus 8B and 70B on the CSCS Alps cluster with 1000+ GH200 nodes.

Reference for infrastructure & compute platforms in the lab *Sep 2023 - Present*
- Support the deployment of container-based research experiments on multiple clusters at the scale of hundreds of nodes, including Slurm HPC clusters and Run:AI (Kubernetes-based) clusters. (Docs.)
- Stack: Docker, Kubernetes, Run:AI, Slurm, NVIDIA Pyxis, NVIDIA Enroot, Apptainer/Singularity.

Open-source contributions and selected projects

Open-source contributions
- @PyTorch/RL, @DLR-RM/stable-baselines3, @Farama-Foundation/Gymnasium, @HuggingFace/TRL.

Python Machine Learning Research Project Template *Sep 2022 - Present*
- Maintain a template for reproducible Python machine-learning research projects deployable on various platforms including Slurm (with containers through NVIDIA Pyxis/Enroot or Apptainer), Kubernetes clusters.

Class projects

SAM - Study Abroad Matching - Python *Sep 2019 – Dec 2019*

- Designed a web app around an algorithm to match students applying to study-abroad programs and universities based on a human-in-the-loop solution to the Gale–Shapley college admissions problem.
- Stack: Django, React, Google Cloud Platform (GCP), Google App Engine, Google Cloud SQL.

Computer facial discrimination - Machine Learning project *May 2019 - Jun 2019*

- Studied the effect of different pre-trained feature extractors (Auto-encoders, CNNs classifiers) and clustering methods on the clustering (computer discrimination) of human faces.
- Built a customizable auto-encoder model to demonstrate that training data as an analogy of background and culture can be the root of multiple kinds of discrimination.
- Stack: Python, TensorFlow, Keras, scikit-learn, Google Colab, Kaggle.

TOP5 Basketball Team Manager - C++ project *Nov 2018 - Jan 2019*

- Managed a team of 9 students / Led the GameEngine team / Maintained the Git repository
- Designed and implemented the game engine and its interaction with the remaining objects of the game
- Stack: C++, Qt.

Meet Halfway - Web project *May 2018 – Jun 2018*

- Built a user-friendly website that allows a group of friends to find the best destination to travel to and meet for holidays according to their respective locations and cost of travel.
- Stack: HTML, CSS, JS.

EDUCATION - EXCHANGE PROGRAMS

University of Toronto **Toronto, Canada**

Exchange - Fall 2019 | Software Engineering, Databases, Compilers & Interpreters *Sep 2019 - Dec 2019*

Stanford University **Stanford, USA**

Exchange - Summer 2019: GPA 4.15/4.0 | AI, Entrepreneurship & Innovation *Jun 2019 - Aug 2019*

LANGUAGES

English: bilingual French: native Arabic: native German: elementary proficiency