

Paul Elvinger

elvingerpa@gmail.com | +352 661 828 192 | elvingerpaul.io | [Google Scholar](#) | [Github](#)

Work Experience

August 2025 – current	Wayve Technologies – Machine Learning Engineer , London <ul style="list-style-type: none">Optimizing training efficiency and performance of Wayve's E2E driving models.
Oct 2024 – June 2025	ETH Zurich – Research Assistant , Zurich <ul style="list-style-type: none">Addressed the underutilization of GPUs by spatially co-locating workloads in an interference-aware manner with a particular focus on ML inference. (<i>Systems Group ETH Zurich, Prof. Ana Klimovic</i>)
Mar 2021 – Apr 2023	Archlet – Data Engineer , Zurich (full time alongside studies) <ul style="list-style-type: none">Part of Archlet's backend team and responsible for the integration, extension and monitoring of several Python based microservices hosted as part of a Kubernetes cluster in Azure.Contributed to the development of Archlet's data pipelines and the operation of its optimization tool backed by mixed integer programming algorithms.
Sep 2020 – Feb 2021	Oracle Labs – Research Assistant , Zurich <ul style="list-style-type: none">Improved the security of Java applications running in containers, using static analysis to automatically generate SELinux policies at application and function granularity. (led to US patent application)

Education

Sep 2022 – Sep 2024	ETH Zurich – MSc. Computer Science (GPA 5.71 / 6.0) <ul style="list-style-type: none">Thesis – Towards resource and interference-aware scheduling of ML workloads on GPUs. [pdf] [code] [slides] (<i>Systems Group ETH Zurich, supervisor: Prof. Ana Klimovic</i>)
Sep 2017 – Sep 2020	ETH Zurich – BSc. Computer Science (GPA: 5.4 / 6.0) <ul style="list-style-type: none">Thesis – Simulated and analyzed the Internet's undergoing process of centralization and its robustness to large scale failures on the physical layer. (<i>Network Security Group ETH Zurich, supervisor: Prof. Adrian Perrig</i>)

Publications and Patents

SoCC'25	Understanding GPU Resource Interference One Level Deeper [pdf] [code] [slides] <u>Paul Elvinger</u> , Foteini Strati, Natalie Enright Jerger, Ana Klimovic, SoCC'25
EuroMLSys '24	Foteini Strati, <u>Paul Elvinger</u> , Tolga Kerimoglu, and Ana Klimovic. [pdf] ML Training with Cloud GPU Shortages: Is Cross-Region the Answer? EuroMLSys 2024.
US Patent 12189759	Neugschwandtner Matthias, Hugo Guioux, and <u>Paul Elvinger</u> . [pdf] Runtime-tailored security policies for applications . U.S. Patent No. 12,189,759. 7 Jan. 2025.

Teaching

Spring 2024	Cloud Computing Architecture , ETH Zurich (Prof. Ana Klimovic)
Fall 2018	Discrete Mathematics , ETH Zurich (Prof. Ueli Maurer)

Programming Skills

Languages	Python, C++/CUDA, Java, SQL
Concepts + Tools	ML Systems, LLMs, PyTorch, NVIDIA Nsight Compute/Systems, GPUs, vLLM, Ray, Kubernetes, Cloud Computing
Cloud Infrastructure	Google Cloud Platform, Microsoft Azure

Languages

English, German, French – Full working proficiency; Luxembourgish – Mother tongue

Extra Curricula

Voluntary Work	Scout Group Leader – Daereldeieren Bridel (FNEL), Luxembourg <ul style="list-style-type: none">Organized weekly meetings for up to 50 children aged from 9 to 13 years.
Sport	Alpine Skiing <ul style="list-style-type: none">Luxembourg National Champion in Slalom 2017Luxembourg Athlete at European Youth Olympic Festival (EYOF) 2015 in Austria
	Golf (Hcp 6.6) <ul style="list-style-type: none">Luxembourg Team Champion 2015, 2016, 2017
	Marathon Runner