



<b>Education</b>	<p>ETH Zurich, Master's Thesis · <i>Machine Learning</i>          TU Wien, Master's Degree · <i>Computer Science</i></p>	<p><i>Jul 2024 - Mar 2025</i>  <i>Oct 2019 - Mar 2025</i></p>
<b>Experience</b>	<p>University of Cambridge  <i>Research Assistant, Part-time</i></p> <ul style="list-style-type: none"> <li>Developing compiler infrastructure for DSLs targeting Machine Learning accelerators (<a href="#">commit history</a>).</li> </ul> <p>Python C++ MLIR LLVM</p> <p>Qube Research &amp; Technologies  <i>Quantitative Analyst (Contract), Full-time</i></p> <ul style="list-style-type: none"> <li>Optimized statistical volatility indicators 2.19x via algorithmic redesign and Rust hot-path migration.</li> <li>Engineered zero-copy pipelines, eliminating memory alignment faults at cross-language FFI boundaries.</li> <li>Built sub-microsecond signal-processing infrastructure for alpha research and production deployment.</li> <li>Modified Zstandard codec internals to support metadata injection with zero performance regression.</li> </ul> <p>Python Rust PyO3 Polars Apache Arrow</p> <p>ETH Zurich  <i>Research Assistant, Full-time</i></p> <ul style="list-style-type: none"> <li>Achieved adversarial robustness at zero inference cost via self-ensembling of intermediate representations.</li> <li>Implemented 16k+ lines of PyTorch reproducing SOTA robustness benchmarks with Google collaborators.</li> <li>Scaled distributed training 5x faster using gradient compression and model parallelism with Lightning.</li> <li>Reduced GPU memory footprint by 40% through large-scale model distillation and memory-efficient kernels.</li> </ul> <p>Python R PyTorch Lightning CUDA Docker</p> <p>University of Waterloo  <i>Research Assistant, Full-time</i></p> <ul style="list-style-type: none"> <li>Isolated bare-metal workloads requiring privileged CPU access using FreeBSD Jails at 2% overhead.</li> <li>Generated synthetic code-embedding datasets via AST permutations and benchmarked dense retrieval.</li> <li>Fast-tracked to PhD admission based on research performance.</li> </ul> <p>Python C PyTorch Faiss Apache Lucene</p> <p>Erste Asset Management  <i>Trading Software Engineer (Contract), Full-time</i></p> <ul style="list-style-type: none"> <li>Designed real-time overlay network for low-latency order synchronization using server-sent events.</li> <li>Cut inference pipeline network overhead 2.5x through request batching and broker-free messaging.</li> <li>Discovered and patched critical timing vulnerability via coordinated server shutdown analysis.</li> </ul> <p>Java Spring-Boot Apache Kafka FIX Protocol</p>	<p>Remote / Cambridge, UK  <i>Dec 2025 - present</i></p> <p>Paris, France  <i>Apr 2025 - Oct 2025</i></p> <p>Zurich, Switzerland  <i>Jul 2024 - Mar 2025</i></p> <p>Ontario, Canada  <i>Sep 2023 - Jun 2024</i></p> <p>Vienna, Austria  <i>Apr 2023 - Sep 2023</i></p>
<b>Publications</b>	<p>Jabary, Y. (2025). Rethinking Adversarial Examples (Master's Thesis, ETH Zurich). (<a href="#">code</a>, <a href="#">pdf</a>)</p> <p>Jabary, Y., Plesner, A., Kuzhagaliyev, T., &amp; Wattenhofer, R. (2024). Seeing Through the Mask: Rethinking Adversarial Examples for CAPTCHAs. (<a href="#">press</a>, <a href="#">code</a>, <a href="#">pdf</a>)</p> <p>Subramanian, V. N., Chang, R., &amp; Jabary, Y. (2023). Benchmarking Vector Databases on Code Embeddings. (<a href="#">pdf</a>)</p>	
<b>Open Source</b>	<p>autograd.c Autograd engine with arena-allocated graph nodes and reference-counted tensor lifetimes. (<a href="#">code</a>)</p> <p>llvm-to-air World's first open end-to-end Apple GPU compiler toolchain through reverse engineering. (<a href="#">code</a>)</p> <p>nogil Outperforming standard library hashing through GIL-bypassing CPython extensions. (<a href="#">code</a>, <a href="#">pdf</a>)</p> <p>Full portfolio at: <a href="https://github.com/sueszli">github.com/sueszli</a></p>	
<b>Awards &amp; Honors</b>	<p>Qube Research &amp; Technologies / Raspberry Pi Foundation, Nominated Contest Judge</p> <p>Google for Developers, Nominated Community Leader</p> <p>EPFL, Summer Research Institute SuRI Fellowship</p> <p>University of Waterloo, Mathematics International Master's Award of Excellence M-IMAE (\$12,500)</p> <p>University of Waterloo, Graduate Research Studentship GRS (\$23,081)</p> <p>TU Wien, Institute of Logic &amp; Computation Research Scholarship (€1,200)</p> <p>Austrian Government, National Maker Competition Prize (€3,000)</p> <p>University of Vienna, Early Admission Program in Mathematics for Gifted Students</p>	<p>2025</p> <p>2024</p> <p>2024</p> <p>2023</p> <p>2023</p> <p>2021</p> <p>2018</p> <p>2017</p>