

James Newling

Location: Santa Cruz, CA, USA
Email: james.newling@gmail.com

Website: <https://newling.github.io>
Github: <https://github.com/newling>

Summary

I am an experienced AI graph compiler engineer, with a strong background in machine learning research. As an enthusiastic and reliable team member, I enjoy taking ownership of the long-term development and success of projects.

Software Skills

Expert in C++ and software design, fluent in ML frameworks, ML compilers, MLIR, OpenCL, and Python, and experienced at developing software for accelerators (AMD GPUs, Graphcore IPU)

Employment

Software Technical Lead (*June 2021 – Present*) at **Graphcore**

I design, implement, and integrate graph compiler passes for lowering from ML frameworks (PyTorch & JAX) to Graphcore's new IPU accelerator. My team uses MLIR for compiler development. I provide technical support and guidance.

Software Team Lead (*May 2020 – June 2021*) at **Graphcore**

I lead a team of 6 engineers developing an IPU-optimized ML framework, PopART.

Member of the Software Team (*March 2018 – May 2020*) at **Graphcore**

I integrate machine learning frameworks into Graphcore's software stack, and implement graph compiler passes. I have designed and implemented algorithms for graph scheduling and outlining, bufferization, liveness and alias analysis, automatic differentiation, model pipelining, parallelization, and serialization, and canonicalizations and constant folding. I collaborate with researchers designing new machine learning models. Some day-to-day technologies are C++, CMake, git, Python, pybind, buildbot, ONNX, PyTorch, LLVM. The role requires a good understanding of the IPU processor, and a focus on efficient use of algorithms and data structures in C++.

External Collaborator (*October 2016 – October 2017*) at **Advanced Micro Devices**

I develop OpenCL GEMM (matrix multiply) kernels for AMD GPUs. This 1 year collaboration is the extension of a PhD internship. The MIOpenGEMM project is a component in AMD's ROCm support for frameworks such as PyTorch.

Maths Lecturer (*February 2010 – June 2010*) at **University of Cape Town**

I am the lecturer for an honours applied maths course in non-linear optimization. Topics covered include artificial neural networks.

Education

February 2018, **PhD in Computer Science** at École Polytechnique Fédérale de Lausanne (EPFL). Supervised by François Fleuret, thesis entitled **Novel Algorithms for Clustering**. I was a member of

the Machine Learning Group at the Idiap Research Institute. The group focused on statistical learning techniques mainly in computer vision, with an interest on computational aspects. My individual research comprised theoretical algorithmic complexity results (which received a best paper award at AISTATS 2017) and optimized open source CPU implementations of popular clustering algorithms.

June 2013, **MSc in Complexity Science** at École Polytechnique (Paris) and Warwick University

June 2011, **Masters in Applied Mathematics** at The University of Cape Town. Masters in observational cosmology, by research.

December 2008, **Honours Degree in Mathematics and Statistics** at The University of Cape Town. Undergraduate degree with 3 years of pure maths and statistics, and courses in applied mathematics, physics, economics. Followed by 1 year in Applied Mathematics for the Honours Degree.

Selected Publications

J. Newling and F. Fleuret. **K-Medoids For K-Means Seeding**. In Proceedings of the International Conference on Neural Information Processing Systems (NIPS), 2017.

J. Newling and F. Fleuret. **A Sub-Quadratic Exact Medoid Algorithm**. In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS), pages 185-193, 2017. *Best paper award*.

J. Newling and F. Fleuret. **Nested Mini-Batch K-Means**. In Proceedings of the International Conference on Neural Information Processing Systems (NIPS), pages 1352-1360, 2016.

J. Newling and F. Fleuret. **Fast K-Means with Accurate Bounds**. In Proceedings of the International Conference on Machine Learning (ICML), pages 936-944, 2016

J. Newling et al. **Statistical classification techniques for photometric supernova typing**. In Monthly Notices of the Royal Astronomical Society. 2011.

Please see <https://github.com/newling> or Google Scholar for a complete list of publications.

Selected University Courses

École Polytechnique Fédérale de Lausanne : Advanced Algorithms, Topics in Theoretical Computer Science, Mathematics of Data, Statistical Physics for Computer Science, Topics on Datacenter Design

Warwick University : Algorithms, Mathematical Biology, Theoretical Neuroscience, Scientific Computing, Fundamentals of Modern Statistical Inference

École Polytechnique : Complex Systems, Dynamical Systems, Numerical ODEs and SDEs, Data Mining, Statistical Learning, Signal Processing, Random Models in Evolution

University of Cape Town : Applied Mathematics (I, II, IV), Computer Science (Ia), Economics (I), Mathematics (I, II, III), Physics (I, II), Statistics (I, II, III)