On the optimization of imaging pipelines

Schoonhoven, R.A.

Bibliography

[1] W. van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabravolski, J. D. Beenhouwer, K. J. Batenburg, and J. Sijbers. "Fast and flexible X-ray tomography using the ASTRA toolbox". *Optics Express* 24.22 (2016), pp. 25129–25147 (cit. on pp. 5, 22, 25, 38).

[2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015 (cit. on p. 8).

[3] S. Abid, F. Fnaiech, and M. Najim. "A new Neural Network pruning method based on the singular value decomposition and the weight initialisation". In: *2002 11th European Signal Processing Conference*. 2002, pp. 1–4 (cit. on p. 61).

[4] S. Akiki, Z. Yang, C. Liu, J. Tang, and S. Liu. "Energy-Aware Automatic Tuning of Many-Core Platform via Gradient Descent". In: *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation*. 2018 (cit. on pp. 107, 117).

[5] J. Ansel, S. Kamil, K. Veeramachaneni, J. Ragan-Kelley, J. Bosboom, U.-M. O'Reilly, and S. Amarasinghe. "OpenTuner: An extensible framework for program autotuning". In: *2014 23rd International Conference on Parallel Architecture and Compilation Techniques (PACT)*. 2014, pp. 303–315 (cit. on pp. 75, 106).

[6] H. Anzt, B. Haugen, J. Kurzak, P. Luszczek, and J. Dongarra. "Experiences in autotuning matrix multiplication for energy minimization on GPUs". *Concurrency and Computation: Practice and Experience* 27.17 (2015) (cit. on pp. 107, 108).

[7] A. H. Ashouri, W. Killian, J. Cavazos, G. Palermo, and C. Silvano. "A Survey on Compiler Autotuning Using Machine Learning". *ACM Computing Surveys* 51.5 (2018) (cit. on pp. 75, 106).

[8] I. P. Astono, J. S. Welsh, S. Chalup, and P. Greer. "Optimisation of 2D U-Net Model Components for Automatic Prostate Segmentation on MRI". *Applied Sciences* 10.7 (2020) (cit. on p. 68).

[9] V. Badrinarayanan, A. Kendall, and R. Cipolla. "Segnet: A deep convolutional encoder-decoder architecture for image segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39.12 (2017), pp. 2481–2495 (cit. on pp. 22, 166).

[10] H. Bal, D. Epema, C. de Laat, R. van Nieuwpoort, J. Romein, F. Seinstra, C. Snoek, and H. Wijshoff. "A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term". *IEEE Computer* 49.5 (May 2016), pp. 54–63 (cit. on pp. 111, 127).

[11] J. Ballé, V. Laparra, and E. P. Simoncelli. "End-to-end optimized image compression". In: *International Conference on Learning Representations (ICLR)*. 2017 (cit. on p. 8).

[12] Bassa, C. G., Romein, J. W., Veenboer, B., van der Vlugt, S., and Wijnholds, S. J. "Fourier-domain dedispersion". *Astronomy & Astrophysics* 657 (2022), A46 (cit. on p. 114).

[13] S. Basu and Y. Bresler. "$O(N^3 \log N)$ backprojection algorithm for the 3-D Radon transform". *IEEE Transactions on Medical Imaging* 21 (2002), pp. 76–88 (cit. on p. 19).

[14] A. Beck and M. Teboulle. "Fast Gradient-Based Algorithms for Constrained Total Variation Image Denoising and Deblurring Problems". *IEEE Transactions on Image Processing* 18.11 (2009), pp. 2419–2434 (cit. on pp. 4, 20, 21).

[15] S. Bhadra, V. A. Kelkar, F. J. Brooks, and M. A. Anastasio. "On hallucinations in tomographic image reconstruction". *IEEE Transactions on Medical Imaging* 40.11 (2021), pp. 3249–3260 (cit. on p. 33).

[16] T. Bicer, D. Gürsoy, R. Kettimuthu, F. D. Carlo, G. Agrawal, and I. T. Foster. "Rapid Tomographic Image Reconstruction via Large-Scale Parallelization". In: *European Conference on Parallel Processing*. 2015 (cit. on p. 19).

[17] C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*. Vol. 4. 4. Springer, 2006 (cit. on p. 6).

[18] D. Blalock, J. J. G. Ortiz, J. Frankle, and J. Guttag. "What is the state of neural network pruning?" *Proceedings of Machine Learning and Systems* (2020) (cit. on pp. 59, 61, 63).

[19] R. Boţ and C. Hendrich. "A Douglas–Rachford type primal-dual method for solving inclusions with mixtures of composite and parallel-sum type monotone operators". *SIAM Journal on Optimization* 24 (Jan. 2013), pp. 2541–2565 (cit. on p. 28).

[20] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018 (cit. on p. 8).

[21] S. Brin and L. Page. "The anatomy of a large-scale hypertextual web search engine". *Computer networks and ISDN systems* 30.1-7 (1998), pp. 107–117 (cit. on pp. 17, 75, 100).

[22] P. C. Broekema, J. J. D. Mol, R. Nijboer, A. van Amesfoort, M. Brentjens, G. M. Loose, W. Klijn, and J. Romein. "Cobalt: A GPU-based correlator and beamformer for LOFAR". *Astronomy and Computing* 23 (2018), pp. 180–192 (cit. on pp. 112, 127).

[23] R. A. Brooks and G. Di Chiro. "Beam hardening in X-ray reconstructive tomography". *Physics in medicine and biology* 21.3 (1976), p. 390 (cit. on p. 47).

[24] G. J. Brostow, J. Fauqueur, and R. Cipolla. "Semantic object classes in video: A high-definition ground truth database". *Pattern Recognition Letters* 30.2 (2009), pp. 88–97 (cit. on pp. 68, 166).

[25] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla. "Segmentation and recognition using structure from motion point clouds". In: *European Conference on Computer Vision (ECCV)*. Springer. 2008, pp. 44–57 (cit. on pp. 68, 166).

[26] M. Burtscher, I. Zecena, and Z. Zong. "Measuring GPU power with the K20 built-in sensor". In: *Proceedings of Workshop on General Purpose Processing Using GPUs*. 2014 (cit. on p. 110).

[27] J.-W. Buurlage, H. Kohr, W. J. Palenstijn, and K. J. Batenburg. "Real-time quasi-3D tomographic reconstruction". *Measurement Science and Technology* 29.6 (2018), p. 064005 (cit. on pp. 19, 22).

[28] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. "A limited memory algorithm for bound constrained optimization". *SIAM Journal on Scientific Computing* 16.5 (1995), pp. 1190–1208 (cit. on p. 82).

[29] E. Calore, S. F. Schifano, and R. Tripiccione. "Energy-performance tradeoffs for HPC applications on low power processors". In: *European Conference on Parallel Processing*. 2015 (cit. on pp. 107, 117).

[30] K. W. Cameron. "Energy Oddities, Part 2: Why Green Computing Is Odd". *IEEE Computer* 46.3 (2013) (cit. on p. 107).

[31] J.-E. Campagne, F. Lanusse, J. Zuntz, A. Boucaud, S. Casas, M. Karamanis, D. Kirkby, D. Lanzieri, Y. Li, and A. Peel. "JAX-COSMO: An End-to-End Differentiable and GPU Accelerated Cosmology Library". *arXiv preprint arXiv:2302.05163* (2023) (cit. on p. 35).

[32] S. Chakraborty, N. K. Nagwani, and L. Dey. "Performance Comparison of Incremental K-means and Incremental DBSCAN Algorithms". *International Journal of Computer Applications* 27.11 (2011), pp. 14–18 (cit. on p. 20).

[33] A. Chaparala, C. Novoa, and A. Qasem. "Autotuning GPU-Accelerated QAP Solvers for Power and Performance". In: *IEEE International Conference on High Performance Computing and Communications, International Symposium on Cyberspace Safety and Security, and International Conference on Embedded Software and Systems*. New York, NY, 2015 (cit. on p. 107).

[34] J. Chen, L. Wu, J. Zhang, L. Zhang, D. Gong, Y. Zhao, Q. Chen, S. Huang, M. Yang, X. Yang, et al. "Deep learning-based model for detecting 2019 novel coronavirus pneumonia on high-resolution computed tomography". *Scientific reports* 10.1 (2020), pp. 1–11 (cit. on p. 33).

[35] T. Chen, B. Xu, C. Zhang, and C. Guestrin. "Training deep nets with sublinear memory cost". *arXiv preprint arXiv:1604.06174* (2016) (cit. on p. 36).

[36] S. Cheng, M. Kim, L. Song, Z. Wu, S. Wang, and N. Hovakimyan. "DiffTune: Auto-Tuning through Auto-Differentiation". *arXiv preprint arXiv:2209.10021* (2022) (cit. on p. 34).

[37] J. W. Choi, D. Bedard, R. Fowler, and R. Vuduc. "A roofline model of energy". In: *IEEE 27th International Symposium on Parallel and Distributed Processing*. 2013 (cit. on p. 108).

[38] S. B. Coban and F. Lucka. *Dynamic 3D X-ray micro-CT data of a tablet dissolution in a water-based gel*. Oct. 2019 (cit. on pp. 28, 68, 167).

[39] S. B. Coban, F. Lucka, W. J. Palenstijn, D. van Loo, and K. J. Batenburg. "Explorative Imaging and Its Implementation At the Flex-Ray Laboratory". *Journal of Imaging* 6.4 (2020), p. 18 (cit. on pp. 28, 40, 68, 167).

[40] T. Connors, A. Qasem, and Q. Yi. "Modeling the Impact of Thread Configuration on Power and Performance of GPUs". *Machine Learning: Theory and Applications* (2015) (cit. on p. 107).

[41] J. Coplin and M. Burtscher. "Effects of source-code optimizations on GPU performance and energy consumption". In: *Proceedings of Workshop on General Purpose Processing Using GPUs*. 2015 (cit. on p. 107).

[42] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*. MIT press, 2009 (cit. on p. 65).

[43] D. R. Danilak. *Why Energy Is A Big And Rapidly Growing Problem For Data Centers*. Forbes Technology Council (cit. on p. 105).

[44] K. Datta, M. Murphy, V. Volkov, S. Williams, J. Carter, L. Oliker, D. Patterson, J. Shalf, and K. Yelick. "Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures". In: *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*. 2008 (cit. on p. 107).

[45] F. De Carlo, D. Gürsoy, D. J. Ching, K. Joost Batenburg, W. Ludwig, L. Mancini, F. Marone, R. Mokso, D. M. Pelt, J. Sijbers, and M. Rivers. "TomoBank: a tomographic data repository for computational X-ray science". *Measurement Science and Technology* 29.3 (2018), p. 034004 (cit. on pp. 14, 45).

[46] E. L. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. "Exploiting linear structure within convolutional networks for efficient evaluation". In: *Advances in Neural Information Processing Systems*. 2014, pp. 1269–1277 (cit. on p. 61).

[47] P Dewdney, W Turner, R Millenaar, R McCool, J Lazio, and T Cornwell. "SKA1 system baseline design". *Document number SKA-TEL-SKO-DD-001 Revision* (2013) (cit. on p. 105).

[48] N. Dhanachandra, K. Manglem, and Y. J. Chanu. "Image segmentation using K-means clustering algorithm and subtractive clustering algorithm". *Procedia Computer Science* 54 (2015), pp. 764–771 (cit. on p. 22).

[49] S. Dittmer, E. J., and P. Maass. "Singular Values for ReLU Layers". *IEEE Transactions on Neural Networks and Learning Systems* (2019), pp. 1–12 (cit. on p. 61).

[50] T. Donath, F. Beckmann, and A. Schreyer. "Automated determination of the center of rotation in tomography data". *Journal of the Optical Society of America A* 23.5 (2006), pp. 1048–1057 (cit. on p. 40).

[51] T. Dong, V. Dobrev, T. Kolev, R. Rieben, S. Tomov, and J. Dongarra. "A step towards energy efficient computing: Redesigning a hydrodynamic application on CPU-GPU". In: *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 2014 (cit. on p. 106).

[52] X. Dong, L. Liu, K. Musial, and B. Gabrys. "NATS-Bench: Benchmarking NAS Algorithms for Architecture Topology and Size". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.7 (2022), pp. 3634–3646 (cit. on p. 87).

[53] X. Dong and Y. Yang. "Network pruning via transformable architecture search". In: *Advances in Neural Information Processing Systems*. 2019, pp. 760–771 (cit. on p. 61).

[54] M. Du, S. Kandel, J. Deng, X. Huang, A. Demortiere, T. T. Nguyen, R. Tucoulou, V. De Andrade, Q. Jin, and C. Jacobsen. "Adorym: A multi-platform generic X-ray image reconstruction framework based on automatic differentiation". *Optics Express* 29.7 (2021), pp. 10000–10035 (cit. on p. 35).

[55] M. Du, Y. S. G. Nashed, S. Kandel, D. Gürsoy, and C. Jacobsen. "Three dimensions, two microscopes, one code: Automatic differentiation for X-ray nanotomography beyond the depth of focus limit". *Science Advances* 6.13 (2020) (cit. on p. 35).

[56] D. Eigen and R. Fergus. "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture". In: *IEEE International Conference on Computer Vision (ICCV)*. 2015, pp. 2650–2658 (cit. on p. 166).

[57] M. Endrizzi. "X-ray phase-contrast imaging". *Nuclear instruments and methods in physics research section A: Accelerators, spectrometers, detectors and associated equipment* 878 (2018), pp. 88–98 (cit. on p. 42).

[58] K. Fan, B. Cosenza, and B. Juurlink. "Accurate Energy and Performance Prediction for Frequency-Scaled GPU Kernels". *Computation* 8.2 (2020) (cit. on pp. 107, 117).

[59] L. A. Feldkamp, L. C. Davis, and J. W. Kress. "Practical cone-beam algorithm". *Journal of the Optical Society of America A* 1.6 (1984), pp. 612–619 (cit. on pp. 4, 20, 37).

[60] J. Filipovič, F. Petrovič, and S. Benkner. "Autotuning of OpenCL kernels with global optimizations". In: *Proceedings of the 1st workshop on autotuning and adaptivity approaches for energy efficient HPC systems*. 2017 (cit. on pp. 76, 106).

[61] J. Filipovič, J. Hozzová, A. Nezarat, J. Ol'ha, and F. Petrovič. "Using hardware performance counters to speed up autotuning convergence on GPUs". *Journal of Parallel and Distributed Computing* 160 (2022), pp. 16–35 (cit. on pp. 76, 106).

[62] D. Freedman, R. Pisani, and R. Purves. *Statistics*. 3rd ed. W.W. Norton, 1998 (cit. on p. 88).

[63] M. Frigo and S. G. Johnson. "The Design and Implementation of FFTW3". *Proceedings of the IEEE* 93.2 (2005), pp. 216–231 (cit. on p. 75).

[64] R. Ge, R. Vogt, J. Majumder, A. Alam, M. Burtscher, and Z. Zong. "Effects of Dynamic Voltage and Frequency Scaling on a K20 GPU". In: *International Conference on Parallel Processing (ICPP)*. 2013 (cit. on pp. 107, 117).

[65] F. Glover. "Tabu search—part I". *ORSA Journal on computing* 1.3 (1989), pp. 190–206 (cit. on p. 79).

[66] G. van Gompel, K. van Slambrouck, M. Defrise, K. J. Batenburg, J. De Mey, J. Sijbers, and J. Nuyts. "Iterative correction of beam hardening artifacts in CT". *Medical physics* 38.S1 (2011), S36–S49 (cit. on pp. 47–49).

[67] R. Goncalves, T. van Tilburg, K. Kyzirakos, et al. "A Spatial Column-store to Triangulate the Netherlands on the Fly." In: *24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. New York, NY, USA: ACM, 2016, 80:1–80:4 (cit. on p. 84).

[68] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. *Deep learning*. Vol. 1. MIT Press, 2016 (cit. on pp. 7, 62).

[69] L. Grady. "Random Walks for Image Segmentation". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28.11 (2006), pp. 1768–1783 (cit. on p. 24).

[70] S. Grauer-Gray, L. Xu, R. Searles, S. Ayalasomayajula, and J. Cavazos. "Auto-tuning a high-level language targeted to GPU codes". In: *Innovative Parallel Computing (InPar)*. 2012, pp. 1–10 (cit. on p. 75).

[71] D. Grewe and A. Lokhmotov. "Automatically Generating and Tuning GPU Code for Sparse Matrix-Vector Multiplication from a High-Level Representation". In: *Proceedings of Workshop on General Purpose Processing Using GPUs*. GPGPU-4. 2011 (cit. on pp. 11, 75, 106).

[72] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008 (cit. on p. 34).

[73] F. C. Groen, I. T. Young, and G. Ligthart. "A comparison of different focus functions for use in autofocus algorithms". *Cytometry: The Journal of the International Society for Analytical Cytology* 6.2 (1985), pp. 81–91 (cit. on p. 40).

[74] J. Guerreiro, A. Ilic, N. Roma, and P. Tomás. "Multi-kernel auto-tuning on GPUs: Performance and energy-aware optimization". In: *23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*. 2015 (cit. on p. 107).

[75] Y. Guo, A. Yao, and Y. Chen. "Dynamic Network Surgery for Efficient DNNs". In: *Advances in Neural Information Processing Systems*. 2016, 1387–1395 (cit. on p. 72).

[76] D. Gürsoy, F. De Carlo, X. Xiao, and C. Jacobsen. "TomoPy: a framework for the analysis of synchrotron tomographic data". *Journal of synchrotron radiation* 21.5 (2014), pp. 1188–1193 (cit. on p. 40).

[77] F. Guzzi, A. Gianoncelli, F. Billè, S. Carrato, and G. Kourousias. "Automatic Differentiation for Inverse Problems in X-ray Imaging and Microscopy". *Life* 13.3 (2023) (cit. on p. 35).

[78] M. P. van Haarlem et al. "LOFAR: The LOw-Frequency ARray". *Astronomy & Astrophysics* 556 (2013) (cit. on pp. 11, 12, 112).

[79] D. Hackenberg, T. Ilsche, J. Schuchart, R. Schöne, W. E. Nagel, M. Simon, and Y. Georgiou. "HDEEM: High Definition Energy Efficiency Monitoring". In: *Energy Efficient Supercomputing Workshop*. 2014, pp. 1–10 (cit. on p. 107).

[80] S. Han, H. Mao, and W. J. Dally. "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding". *International Conference on Learning Representations (ICLR)* (2016) (cit. on p. 59).

[81] S. Han, J. Pool, J. Tran, and W. Dally. "Learning both weights and connections for efficient neural network". In: *Advances in Neural Information Processing Systems*. 2015, pp. 1135–1143 (cit. on p. 62).

[82] B. Hassibi, D. G. Stork, and G. J. Wolff. "Optimal brain surgeon and general network pruning". In: *International conference on neural networks*. IEEE. 1993, pp. 293–299 (cit. on p. 60).

[83] A. B. Hayes, L. Li, D. Chavarría-Miranda, S. L. Song, and E. Z. Zhang. "Orion: A Framework for GPU Occupancy Tuning". In: *17th International Middleware Conference*. 2016 (cit. on p. 107).

[84] K. He, X. Zhang, S. Ren, and J. Sun. "Deep Residual Learning for Image Recognition". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778 (cit. on pp. 24, 64, 68).

[85] Y. He, G. Kang, X. Dong, Y. Fu, and Y. Yang. "Soft Filter Pruning for Accelerating Deep Convolutional Neural Networks". In: *27th International Joint Conference on Artificial Intelligence*. 2018 (cit. on p. 67).

[86] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang. "Filter Pruning via Geometric Median for Deep Convolutional Neural Networks Acceleration". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 (cit. on pp. 61, 67).

[87] Y. He, X. Zhang, and J. Sun. "Channel pruning for accelerating very deep neural networks". In: *IEEE International Conference on Computer Vision (ICCV)*. 2017, pp. 1389–1397 (cit. on p. 60).

[88] S. Heldens, P. Hijma, B. van Werkhoven, et al. "The landscape of exascale research: a data-driven literature analysis". *ACM Computing Surveys* 53.2 (2020), pp. 1–43 (cit. on p. 73).

[89] A. A. Hendriksen. *Mixed-scale Dense Networks for PyTorch*. `https://github.com/ahendriksen/msd_pytorch`. 2020 (cit. on pp. 24, 68).

[90] A. A. Hendriksen, D. Schut, W. J. Palenstijn, N. Viganó, J. Kim, D. M. Pelt, T. van Leeuwen, and K. J. Batenburg. "Tomosipo: fast, flexible, and convenient 3D tomography for complex scanning geometries in Python". *Optics Express* 29.24 (2021), pp. 40494–40513 (cit. on p. 38).

[91] G. T. Herman. "Correction for beam hardening in computed tomography". *Physics in Medicine and Biology* 24.1 (1979), p. 81 (cit. on p. 47).

[92] S. Herrmann. "Determining the difficulty of landscapes by PageRank centrality in local optima networks". In: *Evolutionary Computation in Combinatorial Optimization*. Springer. 2016, pp. 74–87 (cit. on p. 77).

[93] S. Herrmann and F. Rothlauf. "Predicting heuristic search performance with PageRank centrality in local optima networks". In: *Annual Conference on Genetic and Evolutionary Computation*. 2015, pp. 401–408 (cit. on p. 77).

[94] P. Hijma, S. Heldens, A. Sclocco, B. van Werkhoven, and H. E. Bal. "Optimization techniques for GPU programming". *ACM Computing Surveys* 55.11 (2023), pp. 1–81 (cit. on p. 11).

[95] H. H. Holm, A. R. Brodtkorb, and M. L. Sætra. "GPU Computing with Python: Performance, Energy Efficiency and Usability". *Computation* 8.1 (2020) (cit. on p. 107).

[96] R. Horst, P. M. Pardalos, and N. Van Thoai. *Introduction to global optimization*. Springer Science & Business Media, 2000 (cit. on p. 8).

[97] K. Hou, W.-c. Feng, and S. Che. "Auto-tuning strategies for parallelizing sparse matrix-vector (spmv) multiplication on multi-and many-core processors". In: *International Symposium on Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE. 2017, pp. 713–722 (cit. on p. 76).

[98] K.-L. Hua, C.-H. Hsu, S. C. Hidayati, W.-H. Cheng, and Y.-J. Chen. "Computer-aided classification of lung nodules on computed tomography images via deep learning technique". *OncoTargets and therapy* 8 (2015) (cit. on p. 33).

[99] S. Huang, S. Xiao, and W.-c. Feng. "On the energy efficiency of graphics processing units for scientific computing". In: *IEEE 23rd International Symposium on Parallel and Distributed Processing*. 2009 (cit. on p. 107).

[100] Z. Huang, H. Liu, J. Wu, and C. Lv. "Differentiable integrated motion prediction and planning with learnable cost function for autonomous driving". *IEEE Transactions on Neural Networks and Learning Systems* (2023) (cit. on p. 34).

[101] S. Ioffe and C. Szegedy. "Batch normalization: Accelerating deep network training by reducing internal covariate shift". *International Conference on Machine Learning (ICML)* (2015) (cit. on p. 62).

[102] H. Ishibuchi and T. Murata. "A multi-objective genetic local search algorithm and its application to flowshop scheduling". *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 28.3 (1998), pp. 392–403 (cit. on p. 80).

[103] A. K. Jain. *Fundamentals of Digital Image Processing*. USA: Prentice-Hall, Inc., 1989 (cit. on p. 66).

[104] W. Jia, E. Garza, K. A. Shaw, and M. Martonosi. "GPU performance and power tuning using regression trees". *ACM Transactions on Architecture and Code Optimization (TACO)* 12.2 (2015) (cit. on p. 107).

[105] A. S. Jurling and J. R. Fienup. "Applications of algorithmic differentiation to phase retrieval algorithms". *Journal of the Optical Society of America A* 31.7 (2014), pp. 1348–1359 (cit. on p. 35).

[106] A. C. Kak and M. Slaney. *Principles of computerized tomographic imaging*. SIAM, 2001 (cit. on pp. 2, 3, 20, 37).

[107] S. Kamil, C. Chan, L. Oliker, J. Shalf, and S. Williams. "An auto-tuning framework for parallel multicore stencil computations". In: *IEEE International Symposium on Parallel & Distributed Processing (IPDPS)*. 2010 (cit. on pp. 74, 107).

[108] S. Kandel, S. Maddali, M. Allain, S. O. Hruszkewycz, C. Jacobsen, and Y. S. G. Nashed. "Using automatic differentiation as a general framework for ptychographic reconstruction". *Optics Express* 27.13 (2019), pp. 18653–18672 (cit. on p. 35).

[109] E. D. Karnin. "A simple procedure for pruning back-propagation trained neural networks". *IEEE Transactions on Neural Networks* 1.2 (1990), pp. 239–242 (cit. on pp. 59, 60).

[110] J. Kennedy and R. Eberhart. "Particle swarm optimization". In: *International Conference on Neural Networks*. Vol. 4. IEEE. 1995, pp. 1942–1948 (cit. on p. 82).

[111] N. Khouzami, F. Michel, P. Incardona, J. Castrillon, and I. F. Sbalzarini. "Model-based autotuning of discretization methods in numerical simulations of partial differential equations". *Journal of Computational Science* 57 (2022), p. 101489 (cit. on p. 75).

[112] D. Kim, S. Ramani, and J. A. Fessler. "Combining Ordered Subsets and Momentum for Accelerated X-Ray CT Image Reconstruction". *IEEE Transactions on Medical Imaging* 34.1 (2015), pp. 167–178 (cit. on p. 19).

[113] D. Kingma and J. Ba. "Adam: A Method for Stochastic Optimization". *International Conference on Learning Representations (ICLR)* (Dec. 2014) (cit. on pp. 7, 23, 24, 36, 68).

[114] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. "Optimization by simulated annealing". *Science* 220.4598 (1983), pp. 671–680 (cit. on p. 80).

[115] D. Kraft. *A software package for sequential quadratic programming*. Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht. Wiss. Berichtswesen d. DFVLR, 1988 (cit. on p. 82).

[116] A. Krizhevsky, I. Sutskever, and G. E. Hinton. "Imagenet classification with deep convolutional neural networks". *Advances in Neural Information Processing Systems* 25 (2012) (cit. on p. 7).

[117] A. Krzywaniak and P. Czarnul. "Performance/energy aware optimization of parallel applications on GPUs under power capping". In: *International Conference on Parallel Processing and Applied Mathematics*. 2019 (cit. on pp. 107, 111, 117).

[118] M. J. Lagerwerf, W. J. Palenstijn, F. Bleichrodt, and K. J. Batenburg. "An Efficient Interpolation Approach for Exploring the Parameter Space of Regularized Tomography Algorithms". *Fundamenta Informaticae* 172.2 (2020), pp. 143–167 (cit. on pp. 55, 131).

[119] M. J. Lagerwerf, S. B. Coban, and K. J. Batenburg. *High-resolution cone-beam scan of twenty-one walnuts with two dosage levels*. Zenodo. Apr. 2020 (cit. on p. 40).

[120] A. Laugros, R. Schoonhoven, L. Pavlovic, A. Kuan, C. Bosch, A. Hendriksen, A. Diaz, M. Holler, W. C. Lee, A. Schaefer, K. J. Batenburg, P. Cloetens, N. Vigano, and A. Pacureanu. "Self-supervised image restoration in coherent X-ray neuronal microscopy". To be submitted (2024) (cit. on p. 155).

[121] Y. LeCun, Y. Bengio, and G. Hinton. "Deep learning". *Nature* 521.7553 (2015), pp. 436–444 (cit. on pp. 7, 59, 73).

[122] Y. LeCun, J. S. Denker, and S. A. Solla. "Optimal brain damage". In: *Advances in Neural Information Processing Systems*. 1990, pp. 598–605 (cit. on p. 61).

[123] T. van Leeuwen, S. Maretzke, and K. J. Batenburg. "Automatic alignment for three-dimensional tomographic reconstruction". *Inverse Problems* 34.2 (2018), p. 024004 (cit. on p. 38).

[124] C. Li and P. K.-S. Tam. "An iterative algorithm for minimum cross entropy thresholding". *Pattern Recognition Letters* 19 (1998), pp. 771–776 (cit. on pp. 20, 22, 24).

[125] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. "Pruning Filters for Efficient ConvNets". In: *International Conference on Learning Representations (ICLR)*. 2017 (cit. on p. 59).

[126] L. Li and C. Kessler. "MeterPU: a generic measurement abstraction API enabling energy-tuned skeleton backend selection". In: *IEEE Trustcom/BigDataSE/ISPA*. Vol. 3. 2015 (cit. on p. 107).

[127] Y. Li, J. Dongarra, and S. Tomov. "A note on auto-tuning GEMM for GPUs". In: *International Conference on Computational Science (ICCS)*. Springer. 2009, pp. 884–892 (cit. on pp. 11, 74, 75, 106).

[128] R. Lim, B. Norris, and A. Malony. "Autotuning GPU kernels via static and predictive analysis". In: *International Conference on Parallel Processing (ICPP)*. IEEE. 2017, pp. 523–532 (cit. on pp. 76, 106).

[129] C.-S. Lin, S.-M. Teng, and P.-A. Hsiung. "Auto-tuning for GPGPU applications using performance and energy model". *Journal of Systems Architecture* 62 (2016) (cit. on p. 107).

[130] J. Lin, Y. Rao, J. Lu, and J. Zhou. "Runtime neural pruning". In: *Advances in Neural Information Processing Systems*. 2017, pp. 2181–2191 (cit. on pp. 60, 61).

[131] S. Lin, R. Ji, C. Yan, B. Zhang, L. Cao, Q. Ye, F. Huang, and D. Doermann. "Towards Optimal Structured CNN Pruning via Generative Adversarial Learning". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 (cit. on p. 61).

[132] M. Lindauer, K. Eggensperger, M. Feurer, A. Biedenkapp, D. Deng, C. Benjamins, T. Ruhkopf, R. Sass, and F. Hutter. "SMAC3: A versatile Bayesian optimization package for hyperparameter optimization". *The Journal of Machine Learning Research* 23.1 (2022), pp. 2475–2483 (cit. on p. 82).

[133] M. Lindauer, K. Eggensperger, M. Feurer, A. Biedenkapp, D. Deng, C. Benjamins, R. Sass, and F. Hutter. *Sequential Model Algorithm Configuration (SMAC)*. `https://github.com/automl/SMAC3`. 2021 (cit. on p. 83).

[134] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. van der Laak, B. van Ginneken, and C. I. Sánchez. "A survey on deep learning in medical image analysis". *Medical image analysis* 42 (2017), pp. 60–88 (cit. on p. 23).

[135] S. Liu, J. Tang, Z. Zhang, and J.-L. Gaudiot. "Computer architectures for autonomous driving". *Computer* 50.8 (2017), pp. 18–25 (cit. on p. 73).

[136] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi. "A survey of deep neural network architectures and their applications". *Neurocomputing* 234 (2017), pp. 11–26 (cit. on p. 7).

[137] Y. Liu, W. M. Sid-Lakhdar, O. Marques, X. Zhu, C. Meng, J. W. Demmel, and X. S. Li. "GPTune: Multitask learning for autotuning exascale applications". In: *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 2021, 234–246 (cit. on p. 75).

[138] Z. Liu, M. Sun, T. Zhou, G. Huang, and T. Darrell. "Rethinking the value of network pruning". *International Conference on Learning Representations (ICLR)* (2019) (cit. on p. 61).

[139] J. Long, E. Shelhamer, and T. Darrell. "Fully convolutional networks for semantic segmentation". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 3431–3440 (cit. on p. 67).

[140] M. López-Ibáñez, J. Dubois-Lacoste, L. P. Cáceres, M. Birattari, and T. Stützle. "The irace package: Iterated racing for automatic algorithm configuration". *Operations Research Perspectives* 3 (2016), pp. 43–58 (cit. on pp. 82, 87).

[141] H. R. Lourenço, O. C. Martin, and T. Stützle. "Iterated local search". In: *Handbook of metaheuristics*. Springer, 2003, pp. 320–353 (cit. on p. 79).

[142] J.-H. Luo, J. Wu, and W. Lin. "Thinet: A filter level pruning method for deep neural network compression". In: *IEEE International Conference on Computer Vision (ICCV)*. 2017, pp. 5058–5066 (cit. on pp. 59–61).

[143] M. López-Ibáñez, L. Pérez Cáceres, and J. Dubois-Lacoste. *irace: Iterated Racing for Automatic Algorithm Configuration*. `https://github.com/MLopez-Ibanez/irace`. 2021 (cit. on p. 83).

[144] A. K. Maier, C. Syben, B. Stimpel, T. Würfl, M. Hoffmann, F. Schebesch, W. Fu, L. Mill, L. Kling, and S. Christiansen. "Learning with known operators reduces maximum error bounds". *Nature machine intelligence* 1.8 (2019), pp. 373–380 (cit. on p. 35).

[145] A. Mametjanov, D. Lowell, C.-C. Ma, and B. Norris. "Autotuning stencil-based computations on GPUs". In: *IEEE International Conference on Cluster Computing*. 2012, pp. 266–274 (cit. on pp. 75, 106).

[146] C. McLeavy, M. Chunara, R. Gravell, A Rauf, A Cushnie, C. S. Talbot, and R. Hawkins. "The future of CT: deep learning reconstruction". *Clinical radiology* 76.6 (2021), pp. 407–415 (cit. on p. 33).

[147] X. Mei, L. S. Yung, K. Zhao, and X. Chu. "A measurement study of GPU DVFS on energy conservation". In: *Workshop on Power-Aware Computing and Systems*. 2013 (cit. on pp. 107, 117).

[148] C. D. Meyer. *Matrix analysis and applied linear algebra*. Vol. 71. Siam, 2000 (cit. on p. 63).

[149] A. Milesi. *UNet: semantic segmentation with PyTorch*. `https://github.com/milesial/Pytorch-UNet`. 2020 (cit. on p. 68).

[150] Y. Mingqiang, K. Kidiyo, and R. Joseph. "A survey of shape feature extraction techniques". *Pattern Recognition* 15.7 (2008), pp. 43–90 (cit. on p. 20).

[151] M. Mitchell. *An introduction to genetic algorithms*. MIT press, 1998 (cit. on p. 80).

[152] S. Mittal. "A Survey on optimized implementation of deep learning models on the NVIDIA Jetson platform". *Journal of Systems Architecture* 97 (2019), pp. 428–442 (cit. on p. 73).

[153] T. Miyazaki, I. Sato, and N. Shimizu. "Bayesian Optimization of HPC Systems for Energy Efficiency". In: *IEEE International Conference on High Performance Computing*. 2018 (cit. on p. 107).

[154] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz. "Importance Estimation for Neural Network Pruning". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 (cit. on p. 61).

[155] J. J. Moré. "The Levenberg-Marquardt algorithm: Implementation and theory". In: *Numerical Analysis*. Vol. 630. 1978 (cit. on p. 109).

[156] M. C. Mozer and P. Smolensky. "Skeletonization: A Technique for Trimming the Fat from a Network via Relevance Assessment". In: *Advances in Neural Information Processing Systems*. Ed. by D. S. Touretzky. Morgan-Kaufmann, 1989, pp. 107–115 (cit. on pp. 59, 60).

[157] R Muthukrishnan and M. Radha. "Edge detection techniques for image segmentation". *International Journal of Computer Science and Information Technology* 3.6 (2011), p. 259 (cit. on p. 22).

[158] P. Márquez-Neila, L. Baumela, and L. Alvarez. "A Morphological Approach to Curvature-Based Evolution of Curves and Surfaces". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36.1 (2014), pp. 2–17 (cit. on p. 24).

[159] L. Nardi, A. Souza, D. Koeplinger, and K. Olukotun. "HyperMapper: a Practical Design Space Exploration Framework". In: *IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. 2019 (cit. on pp. 75, 107).

[160] Y. S. Nashed, T. Peterka, J. Deng, and C. Jacobsen. "Distributed Automatic Differentiation for Ptychography". *Procedia Computer Science* 108 (2017), pp. 404–414 (cit. on p. 35).

[161] J. A. Nelder and R. Mead. "A simplex method for function minimization". *The computer journal* 7.4 (1965), pp. 308–313 (cit. on p. 81).

[162] Y. E. Nesterov. "A method for solving the convex programming problem with convergence rate $\mathcal{O}(1/k^2)$". In: *Doklady Akademii Nauk*. Vol. 269. Russian Academy of Sciences. 1983, pp. 543–547 (cit. on p. 36).

[163] P. Neubert and P. Protzel. "Compact Watershed and Preemptive SLIC: On Improving Trade-offs of Superpixel Segmentation Algorithms". In: *International Conference on Pattern Recognition*. 2014, pp. 996–1001 (cit. on pp. 20, 22, 24).

[164] J. Nocedal and S. J. Wright. "Conjugate gradient methods". *Numerical optimization* (2006), pp. 101–134 (cit. on p. 81).

[165] C. Nugteren. "CLBlast: A tuned OpenCL BLAS library". In: *International Workshop on OpenCL*. ACM, 2018, 5:1–5:10 (cit. on pp. 84, 112).

[166] C. Nugteren and V. Codreanu. "CLTune: A generic auto-tuner for OpenCL kernels". In: *IEEE International Symposium on Embedded Multicore/ Many-core Systems-on-Chip*. 2015, pp. 195–202 (cit. on pp. 74, 76, 106).

[167] A. Nukada and S. Matsuoka. "Auto-tuning 3-D FFT library for CUDA GPUs". In: *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. ACM. 2009, p. 30 (cit. on p. 74).

[168] NVIDIA. *NVIDIA Management Library (NVML)*. 2011. URL: `https://developer.nvidia.com/nvidia-management-library-nvml` (cit. on p. 110).

[169] G. Ochoa, M. Tomassini, S. Vérel, and C. Darabos. "A study of NK landscapes' basins and local optima networks". In: *Annual Conference on Genetic and Evolutionary Computation*. 2008, pp. 555–562 (cit. on p. 77).

[170] N. Otsu. "A Threshold Selection Method from Gray-Level Histograms". *IEEE Transactions on Systems, Man, and Cybernetics* 9.1 (1979), pp. 62–66 (cit. on pp. 20, 22, 24, 43).

[171] T. Ozturk, M. Talo, E. A. Yildirim, U. B. Baloglu, O. Yildirim, and U. Rajendra Acharya. "Automated detection of COVID-19 cases using deep neural networks with X-ray images". *Computers in Biology and Medicine* 121 (2020), p. 103792 (cit. on p. 33).

[172] D. Paganin, S. C. Mayo, T. E. Gureyev, P. R. Miller, and S. W. Wilkins. "Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object". *Journal of microscopy* 206.1 (2002), pp. 33–40 (cit. on p. 42).

[173] L. Page, S. Brin, R. Motwani, and T. Winograd. *The PageRank citation ranking: Bringing order to the web*. Tech. rep. Stanford InfoLab, 1999 (cit. on pp. 17, 75, 100).

[174] W. Palenstijn, K. Batenburg, and J Sijbers. "Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs)". *Journal of Structural Biology* 176.2 (2011), pp. 250–253 (cit. on p. 19).

[175] G. Pang, C. Yan, C. Shen, A. v. d. Hengel, and X. Bai. "Self-trained deep ordinal regression for end-to-end video anomaly detection". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020, pp. 12173–12182 (cit. on p. 8).

[176] V. Panin, G. Zeng, and G. Gullberg. "Total variation regulated EM algorithm [SPECT reconstruction]". *IEEE Transactions on Nuclear Science* 46.6 (1999), pp. 2202–2210 (cit. on p. 52).

[177] H. Park, Y. W. Ko, J. So, and J.-G. Lee. "Performance/Power Design Space Exploration and Analysis for GPU Based Software". *International Journal of Control and Automation* 6.6 (2013) (cit. on p. 107).

[178] J. Park, S. Li, W. Wen, P. T. P. Tang, H. Li, Y. Chen, and P. Dubey. "Faster CNNs with direct sparse convolutions and guided pruning". *International Conference on Learning Representations (ICLR)* (2017) (cit. on p. 60).

[179] S. Park, S. Latifi, Y. Park, A. Behroozi, B. Jeon, and S. Mahlke. "SRTuner: Effective Compiler Optimization Customization by Exposing Synergistic Relations". In: *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 2022, pp. 118–130 (cit. on p. 75).

[180] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello. "ENet: A deep neural network architecture for real-time semantic segmentation". *International Conference on Learning Representations (ICLR)* (2017) (cit. on pp. 69, 166).

[181] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. "Pytorch: An imperative style, high-performance deep learning library". *Advances in Neural Information Processing Systems* 32 (2019) (cit. on pp. 8, 36).

[182] P. J. Pavan, M. S. Serpa, E. D. Carreño, V. Martínez, E. L. Padoin, P. O. Navaux, J. Panetta, and J.-F. Mehaut. "Improving Performance and Energy Efficiency of Geophysics Applications on GPU Architectures". In: *Latin American High Performance Computing Conference*. 2018 (cit. on pp. 105, 107).

[183] D. M. Pelt and J. A. Sethian. "A mixed-scale dense convolutional neural network for image analysis". *Proceedings of the National Academy of Sciences* 115.2 (2018), pp. 254–259 (cit. on pp. 23, 24, 68, 166).

[184] M Persson, D Bone, and H Elmqvist. "Total variation norm for three-dimensional iterative reconstruction in limited view angle tomography". *Physics in Medicine and Biology* 46.3 (2001), p. 853 (cit. on p. 52).

[185] F. Petrovič, D. Střelák, J. Hozzová, et al. "A benchmark set of highly-efficient CUDA and OpenCL kernels and its dynamic autotuning with Kernel Tuning Toolkit". *Future Generation Computer Systems* 108 (2020), pp. 161–177 (cit. on p. 74).

[186] G. Pilikos, L. Horchens, K. J. Batenburg, T. van Leeuwen, and F. Lucka. "Fast ultrasonic imaging using end-to-end deep learning". In: *IEEE International Ultrasonics Symposium (IUS)*. 2020, pp. 1–4 (cit. on p. 35).

[187] G. Pilikos, C. L. de Korte, T. van Leeuwen, and F. Lucka. "Single Plane-Wave Imaging using Physics-Based Deep Learning". In: *IEEE International Ultrasonics Symposium (IUS)*. 2021, pp. 1–4 (cit. on p. 35).

[188] L. Pineda, T. Fan, M. Monge, S. Venkataraman, P. Sodhi, R. T. Chen, J. Ortiz, D. DeTone, A. Wang, S. Anderson, et al. "Theseus: A library for differentiable nonlinear optimization". *Advances in Neural Information Processing Systems* 35 (2022), pp. 3801–3818 (cit. on p. 34).

[189] M. J. Powell. "A direct search optimization method that models the objective and constraint functions by linear interpolation". In: *Advances in optimization and numerical analysis*. Springer, 1994, pp. 51–67 (cit. on p. 82).

[190] M. J. Powell. "An efficient method for finding the minimum of a function of several variables without calculating derivatives". *The computer journal* 7.2 (1964), pp. 155–162 (cit. on p. 82).

[191] D. C. Price, M. A. Clark, B. R. Barsdell, R. Babich, and L. J. Greenhill. "Optimizing performance-per-watt on GPUs in high performance computing". *Computer Science-Research and Development* 31.4 (2016) (cit. on pp. 107, 109, 117).

[192] G. Procaccianti, P. Lago, A. Vetro, D. M. Fernandez, and R. Wieringa. "The Green Lab: Experimentation in Software Energy Efficiency". In: *IEEE/ACM International Conference on Software Engineering*. 2015 (cit. on p. 107).

[193] M. Puschel, J. M. Moura, J. R. Johnson, et al. "SPIRAL: Code generation for DSP transforms". *Proceedings of the IEEE* 93.2 (2005), pp. 232–275 (cit. on p. 75).

[194] W. Rahmaniar and W.-J. Wang. "Real-Time automated segmentation and classification of calcaneal fractures in CT images". *Applied Sciences* 9.15 (2019), p. 3011 (cit. on p. 20).

[195] A. Rasch, R. Schulze, M. Steuwer, et al. "Efficient Auto-Tuning of Parallel Programs with Interdependent Tuning Parameters via Auto-Tuning Framework (ATF)". *ACM Transactions on Architecture and Code Optimization (TACO)* 18.1 (2021) (cit. on pp. 76, 106, 107).

[196] W. Rawat and Z. Wang. "Deep convolutional neural networks for image classification: A comprehensive review". *Neural computation* 29.9 (2017), pp. 2352–2449 (cit. on p. 7).

[197] D.-Q. Ren and R. Suda. "Global optimization model on power efficiency of GPU and multicore processing element for SIMD computing with CUDA". *Computer Science-Research and Development* 27.4 (2012) (cit. on p. 107).

[198] A. Renda, J. Frankle, and M. Carbin. "Comparing rewinding and fine-tuning in neural network pruning". *International Conference on Learning Representations (ICLR)* (2020) (cit. on pp. 61, 62).

[199] A. Rodriguez, H. Zhang, K. Wiklund, T. Brodin, J. Klaminder, P. Andersson, and M. Andersson. "Refining particle positions using circular symmetry". *PLOS ONE* 12.4 (Apr. 2017), pp. 1–23 (cit. on p. 20).

[200] J. W. Romein and B. Veenboer. "PowerSensor 2: A Fast Power Measurement Tool". In: *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 2018 (cit. on pp. 107, 109).

[201] Romein, John W. "The Tensor-Core Correlator". *Astronomy & Astrophysics* 656 (2021), A52 (cit. on p. 114).

[202] O. Ronneberger, P. Fischer, and T. Brox. "U-net: Convolutional networks for biomedical image segmentation". In: *International Conference on Medical image computing and computer-assisted intervention*. Springer. 2015, pp. 234–241 (cit. on pp. 22, 24, 59, 68).

[203] L. I. Rudin, S. Osher, and E. Fatemi. "Nonlinear total variation based noise removal algorithms". *Physica D: nonlinear phenomena* 60.1-4 (1992), pp. 259–268 (cit. on p. 52).

[204] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. "Learning representations by back-propagating errors". *Nature* 323.6088 (1986), pp. 533–536 (cit. on p. 8).

[205] S. Ryoo, C. I. Rodrigues, S. S. Stone, S. S. Baghsorkhi, S.-Z. Ueng, J. A. Stratton, and W.-m. W. Hwu. "Program optimization space pruning for a multithreaded GPU". In: *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 2008, pp. 195–204 (cit. on pp. 76, 106).

[206] H. Salehinejad and S. Valaee. "Pruning of Convolutional Neural Networks using ising Energy Model". In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2021, pp. 3935–3939 (cit. on p. 60).

[207] E. Saxe. "Power-efficient software". *Communications of the ACM* 53.2 (2010) (cit. on p. 107).

[208] P. Schiffmann, D. Martin, G. Haase, and G. Offner. "Optimizing a RBF Interpolation Solver for Energy on Heterogeneous Systems". In: *Proceedings of the International Conference on Parallel Computing*. Vol. 32. 2017 (cit. on p. 107).

[209] R. Schoonhoven, J. W. Buurlage, D. M. Pelt, and K. J. Batenburg. "Real-time segmentation for tomographic imaging". In: *IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*. 2020, pp. 1–6 (cit. on pp. 68, 155, 167).

[210] R. Schoonhoven, A. A. Hendriksen, D. M. Pelt, and K. J. Batenburg. "Lean: Graph-Based Pruning for Convolutional Neural Networks By Extracting Longest Chains". *CoRR* (2020). arXiv: `2011.06923 [cs.LG]` (cit. on p. 155).

[211] R. Schoonhoven, A. Skorikov, W. J. Palenstijn, D. M. Pelt, A. A. Hendriksen, and K. J. Batenburg. "How auto-differentiation can improve CT workflows: classical algorithms in a modern framework". *Optics Express* 32.6 (2024), pp. 9019–9041 (cit. on p. 155).

[212] R. Schoonhoven, B. Veenboer, B. van Werkhoven, and K. J. Batenburg. "Going green: optimizing GPUs for energy efficiency through model-steered auto-tuning". In: *IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. 2022, pp. 48–59 (cit. on p. 155).

[213] R. Schoonhoven, B. van Werkhoven, and K. J. Batenburg. "Benchmarking optimization algorithms for auto-tuning GPU kernels". *IEEE Transactions on Evolutionary Computation* (2022) (cit. on pp. 117, 155).

[214] R. A. Schoonhoven. *BlooPy: Black-box optimization Python for bitstring, categorical, and numerical discrete problems with local, and population-based algorithms.* `https://github.com/schoonhovenrichard/BlooPy`. 2021 (cit. on p. 83).

[215] R. A. Schoonhoven. *Data and Plotting scripts for GPU Benchmarking 2021 paper.* `https://github.com/schoonhovenrichard/GPU_benchmarking_paper`. 2021 (cit. on p. 87).

[216] R. A. Schoonhoven. *Optimizing CT workflows with auto-differentiation 2023 paper.* `https://github.com/schoonhovenrichard/AutodiffCTWorkflows`. 2023 (cit. on p. 38).

[217] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni. "Green AI". *Communications of the ACM* 63.12 (2020), pp. 54–63 (cit. on p. 105).

[218] A. Sclocco, S. Heldens, and B. van Werkhoven. "AMBER: A real-time pipeline for the detection of single pulse astronomical transients". *SoftwareX* 12 (2020), p. 100549 (cit. on pp. 74, 75).

[219] A. Sclocco. "Accelerating Radio Astronomy with Auto-Tuning" (2017) (cit. on p. 76).

[220] A. Sclocco, H. E. Bal, J. Hessels, J. Van Leeuwen, and R. V. Van Nieuwpoort. "Auto-tuning dedispersion for many-core accelerators". In: *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE. 2014, pp. 952–961 (cit. on pp. 74, 75, 107).

[221] A. Sclocco, J. van Leeuwen, H. E. Bal, and R. V. Van Nieuwpoort. "A real-time radio transient pipeline for ARTS". In: *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. 2015, pp. 468–472 (cit. on pp. 74, 75).

[222] R. Sedgewick and K. Wayne. *Algorithms*. 4th ed. Addison-wesley professional, 2011, pp. 661–666 (cit. on p. 65).

[223] H. Sedghi, V. Gupta, and P. M. Long. "The singular values of convolutional layers". *International Conference on Learning Representations (ICLR)* (2019) (cit. on pp. 61, 66).

[224] E. Y. Sidky, J. H. Jørgensen, and X. Pan. "Convex optimization problem prototyping for image reconstruction in computed tomography with the Chambolle–Pock algorithm". *Physics in Medicine and Biology* 57.10 (2012), p. 3065 (cit. on p. 52).

[225] K. Simonyan and A. Zisserman. "Very deep convolutional networks for large-scale image recognition". *arXiv preprint arXiv:1409.1556* (2014) (cit. on p. 7).

[226] K. Spafford, J. Meredith, and J. Vetter. "Maestro: Data Orchestration and Tuning for OpenCL Devices". In: *European Conference on Parallel Processing*. 2010 (cit. on p. 106).

[227] M. Stachowski, A. Fiebig, and T. Rauber. "Autotuning based on frequency scaling toward energy efficiency of blockchain algorithms on graphics processing units". *Journal of Supercomputing* (2020) (cit. on p. 105).

[228] R. Storn and K. Price. "Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces". *Journal of global optimization* 11.4 (1997), pp. 341–359 (cit. on p. 82).

[229] E. Strubell, A. Ganesh, and A. McCallum. "Energy and Policy Considerations for Deep Learning in NLP". In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2019 (cit. on p. 105).

[230] R. Suda, L. Cheng, and T. Katagiri. "A mathematical method for online autotuning of power and energy consumption with corrected temperature effects". *Procedia Computer Science* 18 (2013) (cit. on p. 107).

[231] M. Thies, F. Wagner, Y. Huang, M. Gu, L. Kling, S. Pechmann, O. Aust, A. Grüneboom, G. Schett, S. Christiansen, and A. Maier. "Calibration by differentiation–Self-supervised calibration for X-ray microscopy using a differentiable cone-beam reconstruction operator". *Journal of Microscopy* 287.2 (2022), pp. 81–92 (cit. on p. 35).

[232] C. Tian, Y. Xu, and W. Zuo. "Image denoising using deep CNN with batch renormalization". *Neural Networks* 121 (2020), pp. 461–473 (cit. on p. 59).

[233] C. Timm, F. Weichert, P. Marwedel, and H. Müller. "Design space exploration towards a realtime and energy-aware GPGPU-based analysis of biosensor data". *Computer Science-Research and Development* 27.4 (2012) (cit. on p. 107).

[234] A. Tiwari, C. Chen, J. Chame, M. Hall, and J. K. Hollingsworth. "A scalable auto-tuning framework for compiler optimization". In: *IEEE International Symposium on Parallel and Distributed Processing (IPDPS)*. 2009, pp. 1–12 (cit. on p. 75).

[235] S. Tomov, R. Nath, H. Ltaief, and J. Dongarra. "Dense linear algebra solvers for multicore with GPU accelerators". In: *International Symposium on Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE. 2010, pp. 1–8 (cit. on pp. 75, 106).

[236] Top500. *List–November 2020*. `https://top500.org/lists/top500/2020/11/`. (accessed November 27, 2020) (cit. on p. 73).

[237] C. Tsallis and D. A. Stariolo. "Generalized simulated annealing". *Physica A* 233.1-2 (1996), pp. 395–406 (cit. on p. 82).

[238] B. Veenboer and J. W. Romein. "Radio-astronomical imaging on graphics processors". *Astronomy and Computing* 32 (2020), p. 100386 (cit. on p. 114).

[239] B. Veenboer and J. W. Romein. "Radio-Astronomical Imaging: FPGAs vs GPUs". In: *European Conference on Parallel Processing*. 2019 (cit. on p. 114).

[240] C. R. Vogel and M. E. Oman. "Iterative Methods for Total Variation Denoising". *SIAM Journal on Scientific Computing* 17.1 (1996), pp. 227–238 (cit. on p. 52).

[241] D. J. Wales and J. P. Doye. "Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms". *The Journal of Physical Chemistry A* 101.28 (1997), pp. 5111–5116 (cit. on p. 81).

[242] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors. "Scikit-image: image processing in Python". *PeerJ* 2 (June 2014), e453 (cit. on p. 28).

[243] Z. Wang, D. Grewe, and M. F. O'boyle. "Automatic and portable mapping of data parallel programs to opencl for gpu-based heterogeneous systems". *ACM Transactions on Architecture and Code Optimization (TACO)* 11.4 (2014), pp. 1–26 (cit. on p. 76).

[244] Z. Wang, X. Xu, N. Xiong, L. T. Yang, and W. Zhao. "Analysis of parallel algorithms for energy conservation with GPU". In: *IEEE/ACM International Conference on Green Computing and Communications & International Conference on Cyber, Physical and Social Computing*. 2010 (cit. on p. 107).

[245] Z. Wang, C. Li, and X. Wang. "Convolutional neural network pruning with structural redundancy reduction". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021, pp. 14913–14922 (cit. on p. 60).

[246] B. van Werkhoven. "Kernel Tuner: A search-optimizing GPU code auto-tuner". *Future Generation Computer Systems* 90 (2018) (cit. on pp. 16, 74, 76, 81, 83, 103, 106, 107, 109).

[247] B. van Werkhoven, J. Maassen, H. E. Bal, et al. "Optimizing convolution operations on GPUs using adaptive tiling". *Future Generation Computer Systems* 30 (2014), pp. 14–26 (cit. on pp. 75, 84).

[248] B. van Werkhoven, W. J. Palenstijn, and A. Sclocco. "Lessons learned in a decade of research software engineering GPU applications". In: *International Conference on Computational Science (ICCS)*. Springer. 2020, pp. 399–412 (cit. on p. 73).

[249] R. C. Whaley and J. J. Dongarra. "Automatically Tuned Linear Algebra Software". In: *Proceedings of the ACM/IEEE Conference on Supercomputing*. 1998 (cit. on p. 75).

[250] F.-J. Willemsen, R. Schoonhoven, J. Filipovič, J. Tørring, R. van Nieuwpoort, and B. van Werkhoven. "A Methodology for Comparing Auto-Tuning Optimization Algorithms". *Future Generation Computer Systems* (2024) (cit. on pp. 130, 132, 155).

[251] P. J. Withers. "X-ray nanotomography". *Materials Today* 10.12 (2007), pp. 26–34 (cit. on p. 42).

[252] A. H. Wright, R. K. Thompson, and J. Zhang. "The Computational Complexity of N-K Fitness Functions". *IEEE Transactions on Evolutionary Computation* 4.4 (2000), 373–379 (cit. on p. 83).

[253] Xizhou Feng, Rong Ge, and K. Cameron. "Power and Energy Profiling of Scientific Applications on Distributed Systems". In: *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. Denver, CO, USA, 2005 (cit. on p. 105).

[254] N. Yang, L. v. Stumberg, R. Wang, and D. Cremers. "D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020, pp. 1281–1292 (cit. on p. 59).

[255] T. Yang, Y. Chen, and V. Sze. "Designing Energy-Efficient Convolutional Neural Networks Using Energy-Aware Pruning". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017, pp. 6071–6079 (cit. on p. 59).

[256] S.-K. Yeom, P. Seegerer, S. Lapuschkin, A. Binder, S. Wiedemann, K.-R. Müller, and W. Samek. "Pruning by explaining: A novel criterion for deep neural network pruning". *Pattern Recognition* 115 (2021), p. 107899 (cit. on p. 61).

[257] M. T. Zeegers, T. van Leeuwen, D. M. Pelt, S. B. Coban, R. van Liere, and K. J. Batenburg. "A tomographic workflow to enable deep learning for X-ray based foreign object detection". *Expert Systems with Applications* 206 (2022), p. 117768 (cit. on p. 49).

[258] M. T. Zeegers. *A collection of 131 CT datasets of pieces of modeling clay containing stones - Part 1 of 5*. Zenodo. Jan. 2022 (cit. on p. 49).

[259] Y. Zhang and F. Mueller. "Auto-generation and auto-tuning of 3D stencil codes on GPU clusters". In: *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 2012, pp. 155–164 (cit. on pp. 11, 75, 106).

[260] Y. Zhang and H. Yu. "Convolutional Neural Network Based Metal Artifact Reduction in X-Ray Computed Tomography". *IEEE Transactions on Medical Imaging* 37.6 (2018), pp. 1370–1381 (cit. on p. 33).

[261] C. Zhao, B. Ni, J. Zhang, Q. Zhao, W. Zhang, and Q. Tian. "Variational Convolutional Neural Network Pruning". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019 (cit. on p. 61).

[262] A. Ziabari, S. Venkatakrishnan, M. Kirka, P. Brackman, R. Dehoff, P. Bingham, and V. Paquit. "Beam Hardening Artifact Reduction in X-Ray CT Reconstruction of 3D Printed Metal Parts Leveraging Deep Learning and CAD Models". In: *ASME International Mechanical Engineering Congress and Exposition*. Vol. 2B: Advanced Manufacturing. American Society of Mechanical Engineers. 2020 (cit. on p. 33).

List of publications

Publications that are part of this thesis:

1. R. Schoonhoven, J. W. Buurlage, D. M. Pelt, and K. J. Batenburg. "Real-time segmentation for tomographic imaging". In: *IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*. 2020, pp. 1–6.

2. R. Schoonhoven, A. Skorikov, W. J. Palenstijn, D. M. Pelt, A. A. Hendriksen, and K. J. Batenburg. "How auto-differentiation can improve CT workflows: classical algorithms in a modern framework". *Optics Express* 32.6 (2024), pp. 9019–9041.

3. R. Schoonhoven, A. A. Hendriksen, D. M. Pelt, and K. J. Batenburg. "Lean: Graph-Based Pruning for Convolutional Neural Networks By Extracting Longest Chains". *CoRR* (2020). arXiv: `2011.06923 [cs.LG]`.

4. R. Schoonhoven, B. van Werkhoven, and K. J. Batenburg. "Benchmarking optimization algorithms for auto-tuning GPU kernels". *IEEE Transactions on Evolutionary Computation* (2022).

5. R. Schoonhoven, B. Veenboer, B. van Werkhoven, and K. J. Batenburg. "Going green: optimizing GPUs for energy efficiency through model-steered auto-tuning". In: *IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. 2022, pp. 48–59.

Publications that are not part of this thesis:

1. F.-J. Willemsen, R. Schoonhoven, J. Filipovič, J. Tørring, R. van Nieuwpoort, and B. van Werkhoven. "A Methodology for Comparing Auto-Tuning Optimization Algorithms". *Future Generation Computer Systems* (2024).

2. A. Laugros, R. Schoonhoven, L. Pavlovic, A. Kuan, C. Bosch, A. Hendriksen, A. Diaz, M. Holler, W. C. Lee, A. Schaefer, K. J. Batenburg, P. Cloetens, N. Vigano, and A. Pacureanu. "Self-supervised image restoration in coherent X-ray neuronal microscopy". To be submitted (2024).