









## REFERENCES

- [1] 2019. AAMAS 2019 tutorial, Scalable deep learning: from theory to practice. <https://sites.google.com/view/scalable-deep-learning>.
- [2] Kale ab Tessera, Sara Hooker, and Benjamin Rosman. 2021. Keep the Gradients Flowing: Using Gradient Flow to Study Sparse Network Optimization. arXiv:2102.01670 [cs.LG]
- [3] Nir Ailon, Omer Leibovich, and Vineet Nair. 2020. Sparse Linear Networks with a Fixed Butterfly Structure: Theory and Practice. arXiv:2007.08864 [cs.LG]
- [4] Zahra Atashgahi, Ghada Sokar, Tim van der Lee, Elena Mocanu, Decebal Constantin Mocanu, Raymond Veldhuis, and Mykola Pechenizkiy. 2020. Quick and Robust Feature Selection: the Strength of Energy-efficient Sparse Training for Autoencoders. arXiv:2012.00560 [cs.LG]
- [5] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. 2018. Deep Rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*.
- [6] Mathijs M. de Weerd, Michael Albert, Vincent Conitzer, and Koos van der Linden. 2018. Complexity of Scheduling Charging in the Smart Grid. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (Stockholm, Sweden) (AAMAS '18)*, 3.
- [7] Tim Dettmers and Luke Zettlemoyer. 2019. Sparse networks from scratch: Faster training without losing performance. arXiv preprint arXiv:1907.04840 (2019).
- [8] S. Dey, K. Huang, P. A. Bearel, and K. M. Chugg. 2019. Pre-Defined Sparse Neural Networks With Hardware Acceleration. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 9, 2 (2019), 332–345.
- [9] Virginia Dignum and Frank Dignum. 2020. Agents Are Dead. Long Live Agents! (AAMAS '20). 1701–1705.
- [10] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. 2020. Rigging the Lottery: Making All Tickets Winners. In *International Conference on Machine Learning*.
- [11] Utku Evci, Yani A. Ioannou, Cem Keskin, and Yann Dauphin. 2020. Gradient Flow in Sparse Neural Networks and How Lottery Tickets Win. arXiv:2010.03533 [cs.LG]
- [12] Jonathan Frankle and Michael Carbin. 2019. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In *International Conference on Learning Representations*.
- [13] Daniel Gebbran, Gregor Verbič, Archie C. Chapman, and Sleiman Mhanna. 2020. Coordination of Prosumer Agents via Distributed Optimal Power Flow: An Edge Computing Hardware Prototype. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS '20)*, 2104–2106.
- [14] Golnaz Habibi, Zachary Kingston, Zijian Wang, Mac Schwager, and James McLurkin. 2015. Pipelined Consensus for Global State Estimation in Multi-Agent Systems (AAMAS '15). 1315–1323.
- [15] Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, 1135–1143.
- [16] Sara Hooker. 2020. The Hardware Lottery. arXiv:2009.06489 [cs.CY]
- [17] Siddhant Jayakumar, Razvan Pascanu, Jack Rae, Simon Osindero, and Erich Elsen. 2020. Top-KAST: Top-K Always Sparse Training. In *Advances In Neural Information Processing Systems*, 1379–1387.
- [18] Norman P Jouppe, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Borden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Computer Architecture (ISCA), 2017 ACM/IEEE 44th Annual International Symposium on*.
- [19] Jeremy Kepner, Simon Alford, Vijay Gadepally, Michael Jones, Lauren Milechin, Ryan Robinett, and Sid Samsi. 2019. Sparse Deep Neural Network Graph Challenge. 2019 *IEEE High Performance Extreme Computing Conference (HPEC)* (2019).
- [20] Jeremy Kepner and Ryan Robinett. 2019. Radix-net: Structured sparse matrices for deep neural networks. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, IEEE, 268–274.
- [21] Koen Kok. 2010. Multi-Agent Coordination in the Electricity Grid, from Concept towards Market Introduction (AAMAS '10). 1681–1688.
- [22] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep Learning. *Nature* 521, 7553 (2015), 436–444.
- [23] Yann LeCun, John S Denker, and Sara A Solla. 1990. Optimal brain damage. In *Advances in neural information processing systems*, 598–605.
- [24] Namhoon Lee, Thalayasingam Ajanthan, and Philip Torr. 2019. SNIP: single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*.
- [25] Junjie Liu, Zhe XU, Runbin SHI, Ray C. C. Cheung, and Hayden K.H. So. 2020. Dynamic Sparse Training: Find Efficient Sparse Network From Scratch With Trainable Masked Layers. In *International Conference on Learning Representations*.
- [26] Shiwei Liu, Decebal Constantin Mocanu, Amarsagar Reddy Ramapuram Matavalam, Yulong Pei, and Mykola Pechenizkiy. 2020. Sparse evolutionary deep learning with over one million artificial neurons on commodity hardware. *Neural Computing and Applications* (2020), 1–16.
- [27] Shiwei Liu, Decebal Constantin Mocanu, Yulong Pei, and Mykola Pechenizkiy. 2021. Selfish Sparse RNN Training. arXiv:2101.09048 [cs.LG]
- [28] Shiwei Liu, Tim van der Lee, Anil Yaman, Zahra Atashgahi, Davide Ferraro, Ghada Sokar, Mykola Pechenizkiy, and Decebal Constantin Mocanu. 2020. Topological Insights into Sparse Neural Networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- [29] Christos Louizos, Max Welling, and Diederik P Kingma. 2018. Learning Sparse Neural Networks through  $L_0$  Regularization. In *International Conference on Learning Representations*.
- [30] Decebal Constantin Mocanu. 2017. *Network computations in artificial intelligence*. Ph.D. Dissertation. Eindhoven University of Technology. <https://pure.tue.nl/ws/portalfiles/portal/69949254>
- [31] Decebal Constantin Mocanu, Elena Mocanu, Phuong H. Nguyen, Madeleine Gibescu, and Antonio Liotta. 2016. A topological insight into restricted Boltzmann machines. *Machine Learning* 104, 2 (01 Sep 2016), 243–270.
- [32] Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H Nguyen, Madeleine Gibescu, and Antonio Liotta. 2018. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications* 9, 1 (2018), 2383.
- [33] E. Mocanu, D. C. Mocanu, P. H. Nguyen, A. Liotta, M. E. Webber, M. Gibescu, and J. G. Slootweg. 2019. On-Line Building Energy Optimization Using Deep Reinforcement Learning. *IEEE Transactions on Smart Grid* 10, 4 (2019), 3698–3708.
- [34] Yousuf Ahmad Mohammad Hasanzadeh Mofrad, Rami Melhem and Mohammad Hammoud. 2021. Accelerating Distributed Inference of Sparse Deep Neural Networks via Mitigating the Straggler Effect. In *In proceedings of IEEE High Performance Extreme Computing (HPEC), Waltham, MA USA*. under review.
- [35] Hesham Mostafa and Xin Wang. 2019. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *International Conference on Machine Learning*.
- [36] Michael C. Mozer and Paul Smolensky. 1989. Using Relevance to Reduce Network Size Automatically. *Connection Science* 1, 1 (1989), 3–16.
- [37] Joshua C. Peterson, David Bourgin, Daniel Reichman, Thomas L. Griffiths, and Stuart J. Russell. 2019. Cognitive model priors for predicting human decisions. In *Proceedings of the 36th International Conference on Machine Learning*.
- [38] Tiago Pinto, Luis Gomes, Pedro Faria, Filipe Sousa, and Zita Vale. 2020. MARTINE: Multi-Agent Based Real-Time Infrastructure for Energy. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS '20)*, 2114–2116.
- [39] Ameya Prabhhu, Girish Varma, and Anoop Nambodiri. 2018. Deep expander networks: Efficient deep networks from graph theory. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 20–35.
- [40] Gobinda Saha and Kaushik Roy. 2021. Gradient Projection Memory for Continual Learning. In *International Conference on Learning Representations*.
- [41] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. 2020. SpaceNet: Make Free Space For Continual Learning. arXiv:2007.07617 [cs.LG]
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, Vol. 30.
- [43] Chaoqi Wang, Guodong Zhang, and Roger Grosse. 2020. Picking Winning Tickets Before Training by Preserving Gradient Flow. In *International Conference on Learning Representations*.
- [44] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Anirudha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. 2020. Supermasks in Superposition. arXiv:2006.14769 [cs.LG]
- [45] Vahid Yazdanpanah, Sara Mehryar, Nicholas R. Jennings, Swenja Surminski, Martin J. Siebert, and Jos van Hillegersberg. 2020. Multiagent Climate Change Research. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS '20)*, 1726–1731.
- [46] X. Zeng, T. Zhi, X. Zhou, Z. Du, Q. Guo, S. Liu, B. Wang, Y. Wen, C. Wang, X. Zhou, L. Li, T. Chen, N. Sun, and Y. Chen. 2020. Addressing Irregularity in Sparse Neural Networks Through a Cooperative Software/Hardware Approach. *IEEE Trans. Comput.* 69, 7 (2020), 968–985.
- [47] S. Zhang, Z. Du, L. Zhang, H. Lan, S. Liu, L. Li, Q. Guo, T. Chen, and Y. Chen. 2016. Cambricon-X: An accelerator for sparse neural networks. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 1–12.
- [48] Shunshi Zhang and Bradley C. Stadie. 2020. One-Shot Pruning of Recurrent Neural Networks by Jacobian Spectrum Evaluation. In *International Conference on Learning Representations*.
- [49] Z. Zhang, D. Zhang, and R. C. Qiu. 2020. Deep reinforcement learning for power system applications: An overview. *CSEE Journal of Power and Energy Systems* 6, 1 (2020), 213–225.
- [50] Aojun Zhou, Yukun Ma, Junnan Zhu, Jianbo Liu, Zhijie Zhang, Kun Yuan, Wenxiu Sun, and Hongsheng Li. 2021. Learning N:M Fine-grained Structured Sparse Neural Networks From Scratch. In *International Conference on Learning Representations*.
- [51] H. Zhu and Y. Jin. 2020. Multi-Objective Evolutionary Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* 31, 4 (2020), 1310–1322.