

Bibliography

- [1] AA Alemi, I Fischer, JV Dillon, and K Murphy. Deep variational information bottleneck. *arXiv e-prints*, 2016. arXiv:1612.00410.
- [2] M AlQuraishi. A watershed moment for protein structure prediction, 2020.
- [3] C Amritkar and Vi Jabade. Image caption generation using deep learning technique. In *Proceedings of the Fourth International Conference on Computing Communication Control and Automation (ICCUBEA)*, pages 1–4. IEEE, 2018.
- [4] C Badue et al. Self-driving cars: A survey. *arXiv e-prints*, 2019. arXiv:1901.04407.
- [5] GA Bello et al. Deep-learning cardiac motion analysis for human survival prediction. *Nature Machine Intell.*, 1(2):95–104, 2019.
- [6] DC Berry and DE Broadbent. Interactive tasks and the implicit-explicit distinction. *British J. of Psychology*, 79(2):251–272, 1988.
- [7] C Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [8] T Chen, S Kornblith, M Norouzi, and G Hinton. A simple framework for contrastive learning of visual representations. *arXiv e-prints*, 2020. arXiv:2002.05709.
- [9] C Doersch. Tutorial on variational autoencoders. *arXiv e-prints*, 2016. arXiv:1606.05908.
- [10] A Esteva, B Kuprel, RA Novoa, J Ko, SM Swetter, HM Blau, and S Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, 2017.
- [11] BH Gary, R Manu, B Tamara, and LM Erik. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report 07-49, University of Massachusetts, Amherst, October 2007.
- [12] P Ghosh, MSM Sajjadi, A Vergari, M Black, and B Schölkopf. From variational to deterministic autoencoders. *arXiv e-prints*, 2019. arXiv:1903.12436.
- [13] IJ Goodfellow, J Pouget-Abadie, M Mirza, B Xu, D Warde-Farley, S Ozair, A Courville, and Y Bengio. Generative adversarial networks. *arXiv e-prints*, 2014. arXiv:1406.2661.

Bibliography

- [14] I Guilliard, R Rogahn, J Piavis, and A Kolobov. Autonomous thermalling as a partially observable Markov decision process. *arXiv e-prints*, 2018. arXiv:1805.09875.
- [15] GE Hinton, TJ Sejnowski, and DH Ackley. Boltzmann machines: Constraint satisfaction networks that learn. Technical report, Dept. Computer Science, Carnegie-Mellon University, 1984.
- [16] M Hutson. From models of galaxies to atoms, simple AI shortcuts speed up simulations by billions of times. *Science*, 367(6479):728, 2020.
- [17] W Jiao et al. A deep learning system can accurately classify primary and metastatic cancers based on patterns of passenger mutations. *BioRxiv*, 2019. doi: 10.1101/214494.
- [18] L Jing and Y Tian. Self-supervised visual feature learning with deep neural networks: A survey. *arXiv e-prints*, 2019. arXiv:1902.06162.
- [19] T Karras, T Aila, S Laine, and J Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. *arXiv e-prints*, 2017. arXiv:1710.10196.
- [20] B Kim, KR Varshney, and A Weller. Proceedings of the 2018 ICML Workshop on Human Interpretability in Machine Learning (WHI 2018). *arXiv e-prints*, 2018. arXiv:1807.01308.
- [21] E Kim, H Kim, K Han, BJ Kang, YM Sohn, OH Woo, and CW Lee. Applying data-driven imaging biomarker in mammography for breast cancer screening: Preliminary study. *Scientific Reports*, 8(1):1–8, 2018.
- [22] DP Kingma, S Mohamed, DJ Rezende, and M Welling. Semi-supervised learning with deep generative models. In *NIPS'14: Proc. 27th International Conference on Neural Information Processing Systems*, pages 3581–3589. MIT Press, 2014.
- [23] DP Kingma and M Welling. Auto-encoding variational Bayes. *arXiv e-prints*, 2013. arXiv:1312.6114.
- [24] CP Langlotz. Will artificial intelligence replace radiologists? *Radiology: Artificial Intelligence*, 1(3), 2019. doi: 10.1148/ryai.2019190058.
- [25] Y LeCun, B Boser, JS Denker, RE Henderson, W Hubbard, and LD Jackel. Backpropagation applied to handwritten ZIP code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [26] G Marcus. Deep learning: A critical appraisal. *arXiv e-prints*, 2018. arXiv:1801.00631.
- [27] MultiMedia LLC. Obama deep fake, 2017.
- [28] I Misra and L van der Maaten. Self-supervised learning of pretext-invariant representations. *arXiv e-prints*, 2019. arXiv:1912.01991.

- [29] V Mnih, K Kavukcuoglu, D Silver, A Graves, I Antonoglou, D Wierstra, and M Riedmiller. Playing Atari with deep reinforcement learning. *arXiv e-prints*, 2013. arXiv:1312.5602.
- [30] V Mnih et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [31] P Noor. Can we trust AI not to further embed racial bias and prejudice? *BMJ*, 368, 2020. doi: 10.1136/bmj.m363.
- [32] E Ntoutsis et al. Bias in data-driven AI systems – an introductory survey. *arXiv e-prints*, 2020. arXiv:2001.09762.
- [33] H Pham, MY Guan, B Zoph, QV Le, and J Dean. Efficient neural architecture search via parameter sharing. *arXiv e-prints*, 2018. arXiv:1802.03268.
- [34] A Radford, L Metz, and S Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv e-prints*, 2015. arXiv:1511.06434.
- [35] DJ Rezende, S Mohamed, and D Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv e-prints*, 2014. arXiv:1401.4082.
- [36] AL Samuel. Some studies in machine learning using the game of checkers. *IBM J. Research and Development*, 3(3):210–229, 1959.
- [37] J Schmidhuber. Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments. Technical report, Institut für Informatik, Technische Universität München, 1990.
- [38] F Schroff, D Kalenichenko, and J Philbin. FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823. IEEE, 2015.
- [39] TJ Sejnowski. *The Deep Learning Revolution*. MIT Press, 2018.
- [40] TJ Sejnowski. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 2020. doi: 10.1073/pnas.1907373117.
- [41] CE Shannon. Programming a computer for playing chess. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(314):256–275, 1950.
- [42] CE Shannon and W Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.
- [43] D Silver et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–503, 2016.
- [44] D Silver et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.

Bibliography

- [45] JD Smith et al. Implicit and explicit categorization: A tale of four species. *Neuroscience & Biobehavioral Reviews*, 36(10):2355–2369, 2012.
- [46] JM Stokes et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702.e13, 2020.
- [47] JV Stone. *Independent Component Analysis: A Tutorial Introduction*. MIT Press, 2004.
- [48] JV Stone. *Bayes' Rule: A Tutorial Introduction to Bayesian Analysis*. Sebtel Press, 2013.
- [49] JV Stone. *Information Theory: A Tutorial Introduction*. Sebtel Press, 2015.
- [50] JV Stone. *Artificial Intelligence Engines: A Tutorial Introduction to the Mathematics of Deep Learning*. Sebtel Press, 2019.
- [51] C Sun, A Shrivastava, S Singh, and A Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)*, pages 843–852. IEEE, 2017.
- [52] RS Sutton and AG Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [53] S Suwajanakorn, SM Seitz, and I Kemelmacher Shlizerman. Synthesizing Obama: Learning lip sync from audio. *ACM Transactions on Graphics*, 36(4):1–13, 2017.
- [54] EJ Topol. High-performance medicine: The convergence of human and artificial intelligence. *Nature Med.*, 25(1):44–56, 2019.
- [55] D Tran, S Cooke, PJ Illingworth, and DK Gardner. Deep learning as a predictive tool for fetal heart pregnancy following time-lapse incubation and blastocyst transfer. *Human Reproduction*, 34(6):1011–1018, 2019.
- [56] X Wang, et al. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2097–2106. IEEE, 2017.
- [57] H Wu and X Gu. Towards dropout training for convolutional neural networks. *Neural Networks*, 71:1–10, 2015.
- [58] Q Yan, DE Weeks, H Xin, A Swaroop, EY Chew, H Huang, Y Ding, and W Chen. Deep-learning-based prediction of late age-related macular degeneration progression. *Nature Machine Intelligence*, 2(2):141–150, 2020.
- [59] AM Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1):1–7, 2019.