Meditations on Artificial Intelligence
Selected Readings

Dan Reed
dan.reed@utah.edu

September 2023

On The Philosophy of Thought

1. Descartes, René, 1637, Discourse on the Method of Rightly Conducting One's Reason and of Seeking Truth in the Science, https://ia801906.us.archive.org/17/items/rmcg0001/Descartes-Discourse-a1.pdf

Computability and Logic

2. Turing, A.M., 2009. Computing Machinery and Intelligence (pp. 23-65). Springer Netherlands. available at https://link.springer.com/chapter/10.1007/978-1-4020-6710-5_3

3. Turing, A. M., On Computable Numbers, with an Application to the Entscheidungsproblem, 1936, available at https://www.cs.virginia.edu/~robins/Turing_Paper_1936.pdf

4. Gödel, K., On Formally Undecidable Propositions of *Principia Mathematica* and Related Systems, 1930, available at https://homepages.uc.edu/~martinj/History_of_Logic/Godel/Godel%20–%20On%20Formally%20Undecidable%20Propositions%20of%20Principia%20Mathematica%201931.pdf

5. Rosenblatt, F., The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain, Psychology Review, 1958, available at http://doi:10.1037/h0042519

6. Sahay, M., Neural Networks and the Universal Approximation Theorem, June 2020, available at https://towardsdatascience.com/neural-networks-and-the-universal-approximation-theorem-8a389a33d30a (see also https://en.wikipedia.org/wiki/Universal_approximation_theorem)

Intelligence and Psychology

7. Hofstadter, D. Gödel, Escher, Bach, Basic Books, 1979 (See: https://en.wikipedia.org/wiki/G%C3%B6del,_Escher,_Bach)

8. Chabris, C. and Simons, D., The Invisible Gorilla: How Our Intuitions Deceive Us, Crown Publishing, 2010, see http://www.theinvisiblegorilla.com

9. Cole, David, "The Chinese Room Argument", The Stanford Encyclopedia of Philosophy (Summer 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.). available a https://plato.stanford.edu/entries/chinese-room/ (see also https://en.wikipedia.org/wiki/Chinese_room)

10. Cardon, D., Cointet, J.P., Mazières, A. and Carey-Libbrecht, L., 2018. Neurons Spike Back. Réseaux, 211(5), pp.173-220. available at https://mazieres.gitlab.io/neurons-spike-back/index.htm

Computing Hardware History and Evolution

11. Burks, A. W., Goldstine, H. H., and von Neumann, J., Preliminary Discussion of the Logical Design of an Electronic Computing Instrument, 1946, available at https://www.ias.edu/sites/default/files/library/Prelim_Disc_Logical_Design.pdf

12. Markoff, J., From PlayStation to Supercomputer for $50,000, NY Times, 2003, available at https://www.nytimes.com/2003/05/26/business/technology-from-playstation-to-supercomputer-for-50000.html

Hardware AI Acceleration

13. NVIDIA Hopper Architecture In-Depth, May 2022, available at https://developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth/

14. Jouppi, N. *et al*, TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings, Proceedings of the International Symposium on Computer Architecture, June 2023, pp. 1-14, available at https://doi.org/10.1145/3579371.3589350

15. AWS Tranium Architecture, available at https://awsdocs-neuron.readthedocs-hosted.com/en/latest/general/arch/neuron-hardware/trainium.html#trainium-arch

16. Reed, D., Gannon, D., Dongarra, J., HPC Forecast: Cloudy and Uncertain, Communications of the ACM, February 2023, Vol. 66 No. 2, pp. 82-90, February 2023, available at https://doi.org/10.1145/3552309

17. Reed, D., Gannon, D., Dongarra, J., Reinventing High Performance Computing: Challenges and Opportunities, March 2022, available at https://arxiv.org/abs/2203.02544

18. Markoff, J., IBM Develops a New Chip that Functions Like the Brain, New York Times, August 7, 2014, available at https://www.nytimes.com/2014/08/08/science/new-computer-chip-is-designed-to-work-like-the-brain.html

19. Ambrogio, S. *et al*, An Analog-AI Chip for Energy-efficient Speech Recognition and Transcription, Nature, 620, pp. 768–775 (2023), available at https://www.nature.com/articles/s41586-023-06337-5

Artificial Intelligence Techniques

20. Parr, T. and Howard, J., The Matrix Calculus You Need For Deep Learning, arXiv preprint, arXiv:1802.01528, July 2018, available at https://arxiv.org/abs/1802.01528

21. GPT-4 Technical Report, March 2023, available at https://arxiv.org/abs/2303.08774

22. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E. and Brynjolfsson, E., 2021. On the Opportunities and Risks of Foundation Models. arXiv preprint arXiv:2108.07258. available at https://arxiv.org/abs/2108.07258

23. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S. and Nori, H., 2023. Sparks of Artificial General Intelligence: Early Experiments with GPT-4, arXiv preprint arXiv:2303.12712. available at https://arxiv.org/abs/2303.12712

24. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D. and Chi, E.H., 2022. Emergent Abilities of Large Language Models. arXiv preprint arXiv:2206.07682. available at https://arxiv.org/abs/2206.07682

25. Narayanan, A., and Kapoor, S., GPT-4 and Professional Benchmarks: The Wrong Answer to the Wrong Question. "AI Snake Oil" blog. March 20, 2023. available at https://aisnakeoil.substack.com/p/gpt-4-and-professional-benchmarks

26. Sevilla, J., Heim, L., Ho, A., Besiroglu, Hobbhahn, M., Villalobos, P., Compute Trends Across Three Eras of Machine Learning, arXiv preprint, 2202.05924v2, February 2022, available at https://arxiv.org/abs/2202.05924

27. Wei, et al, Emergent Abilities of Large Language Models, arXiv preprint, arXiv:2206.07682v2, October 2022, available at https://arxiv.org/abs/2206.07682

28. Goodfellow, I., *et a*l, Generative Adversarial Networks, *Communications of the ACM*, V. 63, No. 11 November 2020, pp 139–144, available at https://doi.org/10.1145/3422622

29. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., 2017. Attention Is All You Need. Advances in Neural Information Processing Systems, 30. 2017, available at https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

30. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. and Agarwal, S., 2020. Language Models are Few-shot Learners. Advances in Neural Information Processing Systems, 33, pp.1877-1901. available at https://arxiv.org/abs/2005.14165

31. Rombach, R., Balttmann, A., Lorenz, D., Esser, Pl, Ommer, B., High-Resolution Image Synthesis with Latent Diffusion Models, arXiv preprint, arXiv:2112.10752v2, December 2021, available at https://doi.org/10.48550/arXiv.2112.10752

32. Li, Y. et al, Textbooks Are All You Need II: phi-1.5 Technical Report, arXiv preprint, arXiv:2309.05463, September 2023, available at https://arxiv.org/abs/2309.05463

Neuroscience Insights for Artificial Intelligence

33. The European Union Human Brain Project, 2023, available at https://www.humanbrainproject.eu/en/

34. White, J. G., Southgate, E., Thomson, J. N., and Brenner, S., The Structure of the Nervous System of the Nematode *Caenorhabditis elegans, Philosophical Transactions of the Royal Society B,* 314: 1–340, 1986, https://royalsocietypublishing.org/doi/10.1098/rstb.1986.0056

35. Winding, M. *et al*, The Connectome of an Insect Brain," *Science*, Vol 379, No. 6636, 2023, https://www.science.org/doi/10.1126/science.add9330

Deep Fakes and Sentiment Analysis

36. This Is Not Morgan Freeman – A Deep Fake Singularity, Diep Nep, 2021, available at https://www.youtube.com/watch?v=oxXpB9pSETo

37. Researchers Use Facial Quirks to Unmask 'Deepfakes,' Neurosciencenews.com, June 2019, available at https://neurosciencenews.com/deepfakes-facial-quirks-14298/

38. Netflix Prize, Wikipedia, available at https://en.wikipedia.org/wiki/Netflix_Prize

Artificial Intelligence Policy, Privacy, and Security

39. Office of Science and Technology Policy, U.S. White House, Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People. October 2022. available at https://www.whitehouse.gov/ostp/ai-bill-of-rights/

40. Feingold, S., The European Union's Artificial Intelligence Act, Explained. World Economic Forum, March 28, 2023. available at https://www.weforum.org/agenda/2023/03/the-european-union-s-ai-act-explained

41. Castelvecchi, D., Can We Open the Black Box of AI?, *Nature* 538, 20–23 (06 October 2016), available at https://doi.org/10.1038/538020a

42. Lawler, R., The US Copyright Office Says You Can't Copyright Midjourney AI-generated Images, The Verge, Feb 22, 2023. available at https://www.theverge.com/2023/2/22/23611278/midjourney-ai-copyright-office-kristina-kashtanova

43. Hill, K., The Technology Facebook and Google Didn't Dare Release, *New York Times*, September 9, 2023, available at https://www.nytimes.com/2023/09/09/technology/google-facebook-facial-recognition.html

Explaining AI Decisions

44. Urbina, F., Lentzos, F., Lnvernizzi, C., Ekins, S., Dual Use of Artificial-Intelligence-Powered Drug Discovery, Nature Machine Intelligence, 4, 198-191, 2022 https://www.nature.com/articles/s42256-022-00465-9

45. Senero, A., and Mezzanotte, M., Open Source Data Science: How to Reduce Bias in AI. Oct. 2022. World Economic Forum. available at https://www.weforum.org/agenda/2022/10/open-source-data-science-bias-more-ethical-ai-technology/

AI Risks and Opportunities

46. Association for the Advancement of Artificial Intelligence. Working Together on Our Future with AI. April 5, 2023. available at https://aaai.org/working-together-on-our-future-with-ai/

47. Bush, V., As We May Think, *Atlantic Monthly*, 1945, available at https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/

Economics, Education, and National Security

48. McKinsey and Company, The Economic Potential of Generative AI, June 2023, available at https://www.mckinsey.com/~/media/mckinsey/business%20functions/mckinsey%20digital/our%20insights/the%20economic%20potential%20of%20generative%20ai%20the%20next%20productivity%20frontier/the-economic-potential-of-generative-ai-the-next-productivity-frontier-vf.pdf

49. Vincent, J., The Lawsuit that Could Rewrite the Rules of AI Copyright., The Verge. November 8, 2022. available at https://www.theverge.com/2022/11/8/23446821/microsoft-openai-github-copilot-class-action-lawsuit-ai-copyright-violation-training-data

50. Plato History, available at http://platohistory.org (see also https://en.wikipedia.org/wiki/PLATO_(computer_system)

AI Winters and Research Funding

51. Minsky, M., and Papert, S., Perceptrons: An Introduction to Computational Geometry,1969. (See also: https://direct.mit.edu/books/book/3132/PerceptronsAn-Introduction-to-Computational)

52. Feigenbaum, E. A., McCorduck, P., The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World, 1983

Scientific Discovery and Innovation

53. Jumper, J., Evans, R., Pritzel, A. *et al.* Highly Accurate Protein Structure Prediction with AlphaFold. *Nature* **596**, 583–589 (2021), available at https://doi.org/10.1038/s41586-021-03819-2

54. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T. and Lillicrap, T., 2018. A General Reinforcement Learning Algorithm that Masters Chess, Shogi, and Go Through Self-play. *Science*, 362(6419), pp.1140-1144. available at https://www.science.org/doi/10.1126/science.aar6404

55. Bi, K., Xie, L., Zhang, H. *et al.* Accurate Medium-range Global Weather Forecasting with 3D Neural Networks. *Nature* **619**, 533–538 (2023), available at https://doi.org/10.1038/s41586-023-06185-3

56. Kuhn, T. S., The Structure of Scientific Revolutions, University of Chicago Press, 1962, (see also https://en.wikipedia.org/wiki/The_Structure_of_Scientific_Revolutions)

57. Guo, S., An Introduction to Surrogate Modeling, Part I: Fundamentals. "Towards Data Science" October 2020. available at https://towardsdatascience.com/an-introduction-to-surrogate-modeling-part-i-fundamentals-84697ce4d241

58. Carter, J., Feddema, J., Kothe, D., Neely, R., Pruet, J., Stevens, R. Advanced Research Directions on AI for Science, Energy, and Security. Eds. Catlett, C. and Dietrich, E., U.S. Department of Energy and National Nuclear Security Administration. May 2023. available at https://www.anl.gov/ai-for-science-report

A Crash Course on Deep Learning

59. Zhou, V., Neural Networks from Scratch, February 2020, available at https://victorzhou.com/series/neural-networks-from-scratch/

60. Malik, M., Basics of Neural Networks, April 2018, available at https://becominghuman.ai/basics-of-neural-network-bef2ba97d2cf

61. Tinker With a Neural Network Right Here in Your Browser, available at http://playground.tensorflow.org

62. Boesch, G. "What's the Difference Between Machine Learning and Deep learning? available at https://viso.ai/deep-learning/deep-learning-vs-machine-learning/

63. Boesch, G., Deep Neural Networks: The 3 Popular Types (MLP, CNN and RNN), Viso.io. available at https://viso.ai/deep-learning/deep-neural-network-three-popular-types/

64. Brownlee, J., A Gentle Introduction to Pooling Layers for Convolutional Neural Networks, Machine Learning Mastery, April 2019. available at https://machinelearningmastery.com/pooling-layers-for-convolutional-neural-networks/

65. Kamali, K. Galaxy Training, Deep Learning (Part 1) - Feedforward Neural Networks (FNN), available at https://training.galaxyproject.org/training-material/topics/statistics/tutorials/FNN/tutorial.html

66. Kamali, K. Galaxy Training, Deep Learning (Part 2) - Recurrent Neural Networks (RNN), available at https://training.galaxyproject.org/training-material/topics/statistics/tutorials/RNN/tutorial.html

67. Olah, C., and Carter, S. Attention and Augmented Recurrent Neural Networks. Distill.pub, September 2016, available at https://distill.pub/2016/augmented-rnns/

68. Zhou, V., Keras for Beginners: Building Your First Neural Network, July 2019, available at https://victorzhou.com/blog/keras-neural-network-tutorial/

69. Alammar, J., The Illustrated Transformer, June 2018, available at http://jalammar.github.io/illustrated-transformer/