Literature

Introduction

https://neilgreenberg.com/ao-quote-borges-on-exactitude-in-science/

Induction & ML

Chapter 1 from Howson, Colin. 2000. Hume's Problem: Induction and the Justification of Belief. 1st ed. Oxford University Press. https://doi.org/10.1093/0198250371.001.0001.

Domingos, Pedro. 2012. "A Few Useful Things to Know about Machine Learning." Communications of the ACM 55 (10): 78–87. https://doi.org/10.1145/2347736.2347755.

Theory

Luxburg, Ulrike von, and Bernhard Schoelkopf. 2008. "Statistical Learning Theory: Models, Concepts, and Results." arXiv:0810.4752. Preprint, arXiv, October 27. https://doi.org/10.48550/arXiv.0810.4752.

Justification

Corfield, David. 2010. "Varieties of Justification in Machine Learning." Minds and Machines 20 (2): 291–301. https://doi.org/10.1007/s11023-010-9191-1.

Decisions

Malik, Momin M. 2020. "A Hierarchy of Limitations in Machine Learning." arXiv:2002.05193. Preprint, arXiv, February 29. https://doi.org/10.48550/arXiv.2002.05193.

Generalization

Lones, Michael A. 2024. "How to Avoid Machine Learning Pitfalls: A Guide for Academic Researchers." Patterns, August, 101046. https://doi.org/10.1016/j.patter.2024.101046.

Recht, Benjamin, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet Classifiers Generalize to ImageNet?

Zhang, Chiyuan, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021. "Understanding Deep Learning (Still) Requires Rethinking Generalization." Communications of the ACM 64 (3): 107–15. https://doi.org/10.1145/3446776.

Lehman, Joel, Elliot Meyerson, Tarek El-Gaaly, Kenneth O. Stanley, and Tarin Ziyaee. 2025. "Evolution and The Knightian Blindspot of Machine Learning." arXiv:2501.13075. Preprint, arXiv, January 22. https://doi.org/10.48550/arXiv.2501.13075.

ML & Statistics

Recht, B. (2024). The Mechanics of Frictionless Reproducibility. Harvard Data Science Review, 6(1). https://doi.org/10.1162/99608f92.f0f013d4)

Breiman, Leo. 2001. “Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author).” Statistical Science 16 (3): 199–231. https://doi.org/10.1214/ss/1009213726.

Chapter 1 and 2 from Shalizi, Cosma Rohilla. 2024. The Truth about Linear Regression. http://www.stat.cmu.edu/~cshalizi/TALR/.

Clustering

Hennig, Christian. 2015. “What Are the True Clusters?” Pattern Recognition Letters 64: 53–62. https://doi.org/10.1016/j.patrec.2015.04.009.

Explanation and Prediction

Shmueli, Galit. 2010. “To Explain or to Predict?” Statistical Science 25 (3): 289–310. https://doi.org/10.1214/10-STS330.

Additional literature (aka you don't need to read this):

Basics

Shalizi, Cosma Rohilla. 2025. Advanced Data Analysis from an Elementary Point of View. http://www.stat.cmu.edu/~cshalizi/ADAfaEPoV/. (This is great but probably 880 pages too long for this seminar.)

Shalizi, Cosma Rohilla. 2022. Course on Data mining https://www.stat.cmu.edu/~cshalizi/dm

Shalizi, Cosma Rohilla. 2022. “Neural Networks.” https://www.stat.cmu.edu/~cshalizi/dm/22/lectures/21/lecture-21.pdf. **(The very basics you need to know about neural networks.)**

Recht, Ben 2025. Rudiments of Prediction. https://people.eecs.berkeley.edu/~brecht/cs281a/prediction_rudiments.pdf (Shows how you can talk about prediction without talking about probability)

General interest

Recht, Benjamin. 2025. “The Actuary’s Final Word on Algorithmic Decision Making.” arXiv:2509.04546. Preprint, arXiv, September 4. https://doi.org/10.48550/arXiv.2509.04546. (What is the place for machines in decision making? What is the place for humans?)

Recht, Ben 2025. Lecture Blog: Pattern, Predictions and Actions. https://www.argmin.net/p/patterns-predictions-and-actions-585

More on generalization

Recht, Ben 2025 "The adaptivity paradox" https://www.argmin.net/p/the-adaptivity-paradox (Why are we not adaptively overfitting?)

Roelofs, Rebecca, Vaishaal Shankar, Benjamin Recht, et al. 2019. “A Meta-Analysis of Overfitting in Machine Learning.” Advances in Neural Information Processing Systems 32. https://proceedings.neurips.cc/paper/2019/hash/ee39e503b6bedf0c98c388b7e8589aca-Abstract.html. (The scientific take on the previous question)

Recht, Ben 2025 "Reshelving generalization" https://www.argmin.net/p/reshelving-generalization (Being pessimistic about generalization theory)

The model complexity myth https://jakevdp.github.io/blog/2015/07/06/model-complexity-myth/#The-Root-of-the-Model-Complexity-Myth (Counting parameters wasn't even a good idea in classical statistics.)

Bengio, Yoshua, and Yves Grandvalet. 2003. “No Unbiased Estimator of the Variance of K-Fold Cross-Validation.” Advances in Neural Information Processing Systems 16. https://proceedings.neurips.cc/paper/2003/hash/e82c4b19b8151ddc25d4d93baf7b908f-Abstract.html. (Does the test error estimate the generalization error?)

Stats & ML

Baum, Eric, and Frank Wilczek. 1987. “Supervised Learning of Probability Distributions by Neural Networks.” In Neural Information Processing Systems, edited by D. Anderson, vol. 0. American Institute of Physics. https://proceedings.neurips.cc/paper_files/paper/1987/file/d8e3068a1b4ad91b2066d3e1780593ee-Paper.pdf. (Apparently they made the first connection between probability distributions and neural nets)

Little, Roderick J A. 2025. Seminal Ideas and Controversies in Statistics. Vol. 176. Monographs on Statistics and Applied Probability. CRC Press. (Chapter 12 contains a discussion of Breiman's paper)

Prediction, Explanation, Model selection

Navarro, Danielle J. 2019. “Between the Devil and the Deep Blue Sea: Tensions Between Scientific Judgement and Statistical Model Selection.” Computational Brain & Behavior 2 (1): 28–34. https://doi.org/10.1007/s42113-018-0019-z.

Giudice, Marco Del. 2024. “The Prediction-Explanation Fallacy: A Pervasive Problem in Scientific Applications of Machine Learning.” Methodology 20 (1): 1. https://doi.org/10.5964/meth.11235.

Uncertainty and related topics:

Hüllermeier, Eyke, and Willem Waegeman. 2021. “Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods.” Machine Learning 110 (3): 457–506. https://doi.org/10.1007/s10994-021-05946-3.

Scorzato, Luigi. 2024. “Reliability and Interpretability in Science and Deep Learning.” Minds and Machines 34 (3): 27. https://doi.org/10.1007/s11023-024-09682-0.

Own the hype:

https://aeon.co/essays/no-suffering-no-death-no-limits-the-nanobots-pipe-dream (A previous hype - how it began, how it ended)