Reading list - Generalization and Overfitting

This reading list is meant to provide some directions through the jungle of statistical and ML literature on the topic of generalization and overfitting. It is designed with the probable discussion topics of the workshop in mind and will hopefully give some background and further reading on them.

Classical definition of generalization and a gentle intro to statistical learning theory:

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

Generalization and simplicity:

Domingos, Pedro. “Occam’s Two Razors: The Sharp and the Blunt,” https://cdn.aaai.org/KDD/1998/KDD98-006.pdf

https://www.yulingyao.com/blog/2023/overfit/

See also the discussion in the comments at:
https://statmodeling.stat.columbia.edu/2023/06/16/bayes-is-guaranteed-to-overfit-what-does-this-mean-theres-a-factor-of-2-here/

Estimating prediction error:

Sections 7.4, 7.5, 7.10 from:
Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. The Elements of Statistical Learning. Springer Series in Statistics. New York, NY: Springer, 2009. https://doi.org/10.1007/978-0-387-84858-7.

Bias-variance trade off:

Section 2.2.2 from:
James, Gareth, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor. An Introduction to Statistical Learning: With Applications in Python. Springer Texts in Statistics. Cham: Springer International Publishing, 2023. https://doi.org/10.1007/978-3-031-38747-0.

Double descent:

Belkin, Mikhail, Daniel Hsu, Siyuan Ma, and Soumik Mandal. “Reconciling Modern Machine-Learning Practice and the Classical Bias–Variance Trade-Off.” Proceedings of the National Academy of Sciences 116, no. 32 (August 6, 2019): 15849–54. https://doi.org/10.1073/pnas.1903070116.

Section 10.8 from
James, Gareth, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor. An Introduction

to Statistical Learning: With Applications in Python. Springer Texts in Statistics. Cham: Springer International Publishing, 2023. https://doi.org/10.1007/978-3-031-38747-0.

http://bactra.org/notebooks/interpolation.html

Classical generalization is not enough:

Zhang, Chiyuan, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. “Understanding Deep Learning (Still) Requires Rethinking Generalization.” Communications of the ACM 64, no. 3 (February 22, 2021): 107–15. https://doi.org/10.1145/3446776.

Maybe it is:
https://workshopping.hlrs.de/2024/01/05/the-illusion-of-generalization/

Sequential overfitting and leakage:

Section 4.3 from:
Lones, Michael A. “How to Avoid Machine Learning Pitfalls: A Guide for Academic Researchers.” arXiv.org, August 5, 2021. https://arxiv.org/abs/2108.02497v4.

Cawley, Gavin C, and Nicola L C Talbot. “On Over-Fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation,” https://www.jmlr.org/papers/volume11/cawley10a/cawley10a.pdf

Kapoor, Sayash, and Arvind Narayanan. “Leakage and the Reproducibility Crisis in Machine-Learning-Based Science.” Patterns 4, no. 9 (September 2023): 100804. https://doi.org/10.1016/j.patter.2023.100804.

Types of inductive justification in ML

Corfield, David. “Varieties of Justification in Machine Learning.” Minds and Machines 20, no. 2 (July 1, 2010): 291–301. https://doi.org/10.1007/s11023-010-9191-1.