

A Survey of Learning Criteria Going beyond the Usual Risk (Abstract Reprint)

Matthew J. Holland¹, Kazuki Tanabe¹

¹Osaka University, Ibaraki, Osaka 567-0047 Japan

Abstract Reprint. This is an abstract reprint of a journal article by Holland and Tanabe (2023).

Abstract

Virtually all machine learning tasks are characterized using some form of loss function, and "good performance" is typically stated in terms of a sufficiently small average loss, taken over the random draw of test data. While optimizing for performance on average is intuitive, convenient to analyze in theory, and easy to implement in practice, such a choice brings about trade-offs. In this work, we survey and introduce a wide variety of non-traditional criteria used to design and evaluate machine learning algorithms, place the classical paradigm within the proper historical context, and propose a view of learning problems which emphasizes the question of "what makes for a desirable loss distribution?" in place of tacit use of the expected loss.

References

Holland, M. J.; and Tanabe, K. 2023. A Survey of Learning Criteria Going Beyond the Usual Risk. *Journal of Artificial Intelligence Research*, 78: 781–821.