Beyond Occam's Razor: Process-Oriented Evaluation

Pedro Domingos

Department of Computer Science and Engineering, University of Washington
Box 352350, Seattle, WA 98195, U.S.A.
pedrod@cs.washington.edu
http://www.cs.washington.edu/homes/pedrod

Abstract. Overfitting is often considered the central problem in machine learning and data mining. When good performance on training data is not enough to reliably predict good generalization, researchers and practitioners often invoke "Occam's razor" to select among hypotheses: prefer the simplest hypothesis consistent with the data. Occam's razor has a long history in science, but a mass of recent evidence suggests that in most cases it is outperformed by methods that deliberately produce more complex models. The poor performance of Occam's razor can be largely traced to its failure to account for the search process by which hypotheses are obtained: by effectively assuming that the hypothesis space is exhaustively searched, complexity-based methods tend to over-penalize large spaces. This talk describes how information about the search process can be taken into account when evaluating hypotheses. The expected generalization error of a hypothesis is computed as a function of the search steps leading to it. Two variations of this "process-oriented" approach have yielded significant improvements in the accuracy of a rule learner. Process-oriented evaluation leads to the seemingly paradoxical conclusion that the same hypothesis will have different expected generalization errors depending on how it was generated. I believe that this is as it should be, and that a corresponding shift in our way of thinking about inductive learning is required.

References

1. P. Domingos. Occam's two razors: The sharp and the blunt. *Proc. 4th Intl. Conf. on Knowledge Discovery and Data Mining*, pp. 37–43, New York, NY, 1998. AAAI Press.
2. P. Domingos. A process-oriented heuristic for model selection. *Proc. 15th Intl. Conf. on Machine Learning*, pp. 127–135, Madison, WI, 1998. Morgan Kaufmann.
3. P. Domingos. Process-oriented estimation of generalization error. *Proc. 16th Intl. Joint Conf. on Artificial Intelligence*, pp. 714–719, Stockholm, Sweden, 1999. Morgan Kaufmann.
4. P. Domingos. The role of Occam's razor in knowledge discovery. *Data Mining and Knowledge Discovery*, 3:409–425, 1999.

R. López de Mántaras, E. Plaza (Eds.): ECML 2000, LNAI 1810, p. 3, 2000.