

The evolution of learning systems: to Bayes or not to be

Nestor Caticha

Instituto de Física Universidade de São Paulo,
Caixa Postal 66318 CEP 05315-970 São Paulo, Brazil

Abstract

Bayesian algorithms pose a limit to the performance learning algorithms can achieve. Natural selection should guide the evolution of information processing systems towards those limits. What can we learn from this evolution and what properties do the intermediate stages have? While this question is too general to permit any answer, progress can be made by restricting the class of information processing systems under study. We present analytical and numerical results for the evolution of on-line algorithms for learning from examples for neural network classifiers, which might include or not a hidden layer. The analytical results are obtained by solving a variational problem to determine the learning algorithm that leads to maximum generalization ability. Simulations using evolutionary programming, for programs that implement learning algorithms, confirm and expand the results.

The principal result is not just that the evolution is towards a Bayesian limit and that indeed it is essentially reached. In addition we find that evolution is driven by the discovery of useful structures or combinations of variables and operators. In different runs the temporal order of the discovery of such combinations is unique. The main result is that combinations that signal the surprise brought by an example arise always before combinations that serve to gauge the performance of the learning algorithm. This latter structures can be used to implement annealing schedules. The temporal ordering can be understood analytically as well by doing the functional optimization in restricted functional spaces. We also show that there is data suggesting that the appearance of these traits also follows the same temporal ordering in biological systems.