Contents

Bayesian Back-Propagation

Wray L. Buntine
wray@ptolemy.arc.nasa.gov
RIACS & NASA Ames Research Center
Mail Stop 269-2, Moffet Field, CA 94035, USA

Andreas S. Weigend
andreas@psych.stanford.edu
Xerox Palo Alto Research Center
3333 Coyote Hill Rd., Palo Alto, CA, 94304, USA

November 21, 1991

Abstract

Connectionist feed-forward networks, trained with back-propagation, can be used both for non-linear regression and for (discrete one-of-C) classification. This paper presents approximate Bayesian methods to statistical components of back-propagation: choosing a cost function and penalty term (interpreted as a form of prior probability), pruning insignificant weights, estimating the uncertainty of weights, predicting for new patterns ("out-of-sample"), estimating the uncertainty in the choice of this prediction ("error bars"), estimating the generalization error, comparing different network structures, and handling missing values in the training patterns. These methods extend some heuristic techniques suggested in the literature, and in most cases require a small additional factor in computation during back-propagation, or computation once back-propagation has finished.

To appear in December 1991 issue of *Complex Systems*. An early draft appeared in the Neuroprose archive July 1991, and was available as NASA Ames AI Research Branch TR FIA-91-22.

This final version contains many corrections and additions, so I was loath to keep the early version available. My apologies for not including the final version in full in the Neuroprose archive, however I felt it would be unfair to the journal. Andreas is on holidays so blame and flames to me. — Wray.