

Neural Network Guided Search Control in Partial Order Planning

Terry Zimmerman

Advisor: Subbarao Kambhampati

Department of Computer Science and Engineering
Arizona State University, Tempe AZ 85287-5406
{zim@asu.edu}

The development of efficient search control methods is an active research topic in the field of planning (Kambhampati, Katukam, & Qu 1996). Investigation of a planning program integrated with a neural network (NN) that assists in search control is underway, and has produced promising preliminary results.

Project Overview: The UCPOP partial order planner (Penberthy & Weld 1992) was used in this project and the initial experiments were limited to "blocks world" problems with up to 3 blocks. Experimentation was done with several candidate sets of "partial plan" parameters or metrics in the search for one that is useful to the NN in discerning whether a partial plan is likely to evolve into a solution plan. This parameter set functions as the input vector to the NN.

The planning program was modified to automatically produce an input vector for every partial plan visited in its search process. When search is complete the program classifies each vector associated with plans lying on the solution path as positive examples and all other vectors generated as negative examples. The modified planner is run on a representative set of problems, generating two sets of input vectors - one for network training, one for testing.

The next project phase involved designing and training a NN using the training set of input vectors. The network should correctly classify a maximum number of its training vectors and still be able to generalize over vectors not yet presented. The trained network is subsequently tested on the set of input vectors **not** used in training. The input vector and/or network design is revisited if the trained network does not perform well in classifying these partial plan vectors.

The final phase is to develop a version of the UCPOP that incorporates the threshold function representing the successfully trained NN in such a manner that it can efficiently guide the planner's solution search. The performance of the modified planner on a variety of problems can then be compared with the original planner.

Current Design: The modified UCPOP program generated roughly 500 input vectors for training from 6 problems and 1000 testing vectors from a different 13 problems. UCPOP's default Best First Search (BFS) algorithm was used for search control during the input vector generation phase.

Designing a "good" candidate set of input parameters for the NN is a major (and ongoing) issue requiring careful analysis and iteration. But the process in itself provides interesting insight into the open question of the "goodness" of a partial plan. The candidate set on which the most

extensive experimentation was performed contained 27 partial plan parameters:

- Number of nodes expanded at the current search stage
- Number of "refinements" in the current plan
- Number of unsatisfied subgoals on the goals queue
- 12 parameters that can take on 1 of 3 values representing whether each of the possible states for blocks (such as on(A B)) are present in the initial and/or goal state
- 12 related parameters indicating whether a series of (possibly temporary) causal links to the initial state has been established for each goal state literal

In this case the best performance was obtained with a multilayer feedforward network with 20 hidden nodes and one output node. The network was trained using a variation of the Gauss-Newton algorithm (Haykin 1994).

Results and Current Directions: When presented with the 1000 vectors from problems not trained on the NN correctly classifies 82% of them. That is, the trained NN exhibits a success rate of 82% in discriminating between those partial plans which will ultimately lie on the solution path and those which won't for plans created and visited by the planner during its solution search for test problems *not previously "seen" by the neural net*. This suggests the NN as configured has a "forecasting" capability that may be useful in improving the planner's search control.

Work is currently still underway on the last project phase discussed above, but initial attempts to use the "NN forecaster" in combination with UCPOP's BFS algorithm to provide search control have produced mixed results. It is apparent that UCPOP performance is very sensitive to how the NN forecasting is used to guide its search control. Various techniques for combining the BFS and NN guidance as well as methods for using the NN guidance alone are being investigated. There are also a number of interesting possibilities and techniques for improving the input vector with respect to partial plan parameters likely to be most useful to the neural network forecasting process.

References

- Haykin, S. 1994. *Neural Networks - A Comprehensive Foundation*. Macmillan College Publishing, Inc. 215--217.
- Kambhampati, S.; Katukam, S.; and Qu, Y. 1996. Failure driven dynamic search control for partial order planners: An explanation based approach. *Artificial Intelligence (forthcoming)*.
- Penberthy, J., and Weld, D. 1992. Ucpop: A sound, complete, partial order planner for adl. *Proceedings of KR-92* 103--114.