Pseudo-Boolean Competition 2010 Solver Description: `borg-pb-10.05.30`

Bryan Silverthorn and **Risto Miikkulainen**
Department of Computer Science
The University of Texas at Austin
{bsilvert,risto}@cs.utexas.edu

In Brief

Algorithm portfolio methods (Huberman, Lukose, and Hogg 1997) use information about solvers and problem instances to allocate computational resources among multiple solvers, attempting to maximize the time spent on those well suited to each instance. Portfolio methods such as SATzilla (Xu et al. 2008) have proved increasingly effective in satisfiability.

An algorithm portfolio must decide which solvers to run and for how long to run them. These decisions rely entirely on expectations about solver behavior.

The `borg` solver attempts to to learn predictable aspects of solver behavior—such as how likely a solver is to succeed if it has previously failed—given data on the successes and failures of solvers on many problem instances. The version of this solver submitted to the 2010 pseudo-Boolean competition, `borg-pb-10.05.30`, assumes a specific *latent class* model of solver behavior, a mixture of Dirichlet compound multinomial (DCM) distributions, which is used to identify groups of similar problem instances. This model is examined in detail by Silverthorn and Miikkulainen (2010). It captures the basic correlations between solvers, runs, and problem instances, as well as the tendency of solver outcomes to recur. Unlike the classifier employed by SATzilla, the model considers only the success or failure of each past solver run; it does *not* consider instance feature information.

This version of `borg-pb` combines the DCM mixture model with a deterministic greedy action selection policy, which yields a fixed execution schedule followed for every problem instance. This policy choice is also discussed further by Silverthorn and Miikulainen (2010).

Portfolio Composition

Portfolio methods rely entirely on the performance of the solvers they employ, and are possible only because of the engineering and research involved in making those solvers effective. This version of `borg-pb` considered 8 subsolvers in its model, several groups of which were multiple parameterizations of the same solver. Table 1 lists these solvers and their authors. Note that the final policy did not necessarily use every solver.

Acknowledgments

This research was supported in part by the NSF under grants EIA-0303609, IIS-0757479, and IIS-0915038, and by the THECB under grant 003658-0036-2007.

References

Huberman, B.; Lukose, R.; and Hogg, T. 1997. Economics Approach to Hard Computational Problems. *Science*.

Silverthorn, B., and Miikkulainen, R. 2010. Latent Class Models for Algorithm Portfolio Methods. In *AAAI*.

Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. SATzilla: Portfolio-based Algorithm Selection for SAT. *JAIR*.

Name	Author(s)
clasp-1.3.3	Martin Gebser, Benjamin Kaufmann, and Torsten Schaub
sat4j-pb-v20090829	Daniel Le Berre and Anne Parrain
sat4j-pb-v20090829-cutting	Daniel Le Berre and Anne Parrain
bsolo_pb10-l1	Vasco Manquinho
bsolo_pb10-l2	Vasco Manquinho
bsolo_pb10-l3	Vasco Manquinho
wbo1.4a	Vasco Manquinho
wbo1.4b	Vasco Manquinho

Table 1: Subsolvers considered by borg-pb-10.05.30.