

## Bayesian Optimisation for Nurse Scheduling

Uwe Aickelin and Jingpeng Li (speaker)

{uxa, jpl}@cs.nott.ac.uk

School of Computer Science and IT

University of Nottingham

Nottingham, NG8 1BB, UK

Combinatorial Optimization Conference (CO2004), Lancaster, UK, 2004.

A Bayesian optimisation algorithm for a nurse scheduling problem is presented, which involves choosing a suitable scheduling rule from a set for each nurse's assignment. When a human scheduler works, he normally builds a schedule systematically following a set of rules. After much practice, the scheduler gradually masters the knowledge of which solution parts go well with others. He can identify good parts and is aware of the solution quality even if the scheduling process is not yet completed, thus having the ability to finish a schedule by using flexible, rather than fixed, rules. In this paper, we design a more human-like scheduling algorithm, by using a Bayesian optimisation algorithm to implement explicit learning from past solutions. A nurse scheduling problem from a UK hospital is used for testing.

Unlike our previous work that used Genetic Algorithms to implement implicit learning [1], the learning in the proposed algorithm is explicit, i.e. we identify and mix building blocks directly. The Bayesian optimisation algorithm is applied to implement such explicit learning by building a Bayesian network of the joint distribution of solutions. The conditional probability of each variable in the network is computed according to an initial set of promising solutions. Subsequently, each new instance for each variable is generated by using the corresponding conditional probabilities, until all variables have been generated, i.e. in our case, new rule strings have been obtained. Sets of rule strings are generated in this way, some of which will replace previous strings based on fitness. If stopping conditions are not met, the conditional probabilities for all nodes in the Bayesian network are updated again using the current set of promising rule strings.

For clarity, consider the following toy example of scheduling five nurses with two rules (1: random allocation, 2: allocate nurse to low-cost shifts). In the beginning of the search, the probabilities of choosing rule 1 or 2 for each nurse is equal, i.e. 50%. After a few iterations, due to the selection pressure and reinforcement learning, we experience two solution pathways: Because pure low-cost or random allocation produces low quality solutions, either rule 1 is used for the first 2-3 nurses and rule 2 on remainder or vice versa. In essence, Bayesian network learns 'use rule 2 after 2-3x using rule 1' or vice versa.

It should be noted that for our and most other scheduling problems, the structure of the network model is known and all variables are fully observed. In this case, the goal of learning is to find the rule values that maximize the likelihood of the training data. Thus, learning can amount to 'counting' in the case of multinomial distributions. For our problem, we use our rules: Random, Cheapest Cost, Best Cover and Balance of Cost and Cover.

In more detail, the steps of our Bayesian optimisation algorithm for nurse scheduling are:

1. Set  $t = 0$ , and generate an initial population  $P(0)$  at random;
2. Use roulette-wheel selection to choose a set of promising rule strings  $S(t)$  from  $P(t)$ ;
3. Compute conditional probabilities of each node according to this set of promising solutions;

4. Assign each nurse using roulette-wheel selection based on the rules' conditional probabilities. A set of new rule strings  $O(t)$  will be generated in this way;
5. Create a new population  $P(t+1)$  by replacing some rule strings from  $P(t)$  with  $O(t)$ , and set  $t = t+1$ ;
6. If the termination conditions are not met (we use 2000 generations), go to step 2.

Computational results from 52 real data instances demonstrate the success of this approach. They also suggest that the learning mechanism in the proposed approach might be suitable for other scheduling problems. Another direction for further research is to see if there is a good constructing sequence for individual data instances, given a fixed nurse scheduling order. If so, the good patterns could be recognized and then extracted as new domain knowledge. Thus, by using this extracted knowledge, we can assign specific rules to the corresponding nurses beforehand, and only schedule the remaining nurses with all available rules, making it possible to reduce the solution space.

#### Acknowledgements

The work was funded by the UK Government's major funding agency, Engineering and Physical Sciences Research Council (EPSRC), under grant GR/R92899/01.

#### References

- [1] Aickelin U, "An Indirect Genetic Algorithm for Set Covering Problems", *Journal of the Operational Research Society*, 53(10): 1118-1126,