

MAP assignment

structure search

This approach can be applied to a broad range of problems. For example, we can use it to find a *MAP assignment* relative to a distribution P : the space of solutions is the set of assignments ξ to a set of random variables \mathcal{X} ; the objective function is $P(\xi)$; and the search operators take one assignment x and change the value of one variable X_i from x_i to x'_i . As we discuss in section 18.4, it can also be used to perform *structure search* over the space of Bayesian network structures to find one that optimizes a certain “goodness” function: the search space is the set of network structures, and the search operators make small changes to the current structure, such as adding or deleting an edge.

Algorithm A.5 Greedy local search algorithm with search operators

```

Procedure Greedy-Local-Search (
     $\sigma_0$ , // initial candidate solution
    score, // Score function
     $\mathcal{O}$ , // Set of search operators
)
1   $\sigma_{\text{best}} \leftarrow \sigma_0$ 
2  do
3       $\sigma \leftarrow \sigma_{\text{best}}$ 
4      Progress  $\leftarrow$  false
5      for each operator  $o \in \mathcal{O}$ 
6           $\sigma_o \leftarrow o(\sigma)$  // Result of applying  $o$  on  $\sigma$ 
7          if  $\sigma_o$  is legal solution then
8              if score( $\sigma_o$ ) > score( $\sigma_{\text{best}}$ ) then
9                   $\sigma_{\text{best}} \leftarrow \sigma_o$ 
10                 Progress  $\leftarrow$  true
11 while Progress
12
13 return  $\sigma_{\text{best}}$ 

```

A.4.2.1 Local Hill Climbing
greedy
hill-climbing

One of the simplest, and often used, search procedures is the *greedy hill-climbing* procedure. As the name suggests, at each step we take the step that leads to the largest improvement in the score. This is the search analogue of a continuous gradient-ascent method; see appendix A.5.2. The actual details of the procedure are shown in algorithm A.5. We initialize the search with some solution σ_0 . Then we repeatedly execute the following steps: We consider all of the solutions that are neighbors of the current one, and we compute their score. We then select the neighbor that leads to the best improvement in the score. We continue this process until no modification improves the score. One issue with this algorithm is that the number of operators that can be applied may be quite large. A slight variant of this algorithm, called *first-ascent hill climbing*, samples operators from \mathcal{O} and evaluates them one at a time. Once it finds one that leads to better scoring network, it applies it without considering other operators. In the initial stages of the search, this procedure requires relatively few random trials before it finds such an

first-ascent hill
climbing