



國立成功大學  
National Cheng Kung University

# Introduction to Artificial Intelligence

## Chapter 4 Search in Complex Environments

Wei-Ta Chu (朱威達)


# Local Search Algorithms and Optimization Algorithms

- The search algorithms that we have seen so far are designed to explore search spaces systematically. When a goal is found, the path to that goal also constitutes a solution to the problem.
- In many problems, however, the path to the goal is irrelevant. In the 8-queens problem, what matters is the final configuration of queens, not the order in which they are added.
  - We need algorithms not worrying about paths at all.

# Local Search Algorithms and Optimization Algorithms

- **Local search** algorithms operate using a single current node and generally move only to neighbors of that node.
  - They use very little memory—usually a constant amount
  - They can often find reasonable solutions in large or infinite (continuous) state spaces for which systematic algorithms are unsuitable.
- Local search algorithms are useful for solving pure **optimization problems**, in which the aim is to find the best state according to an **objective function**.

# Local Search Algorithms and Optimization Algorithms



**Figure 4.1** A one-dimensional state-space landscape in which elevation corresponds to the objective function. The aim is to find the global maximum. Hill-climbing search modifies the current state to try to improve it, as shown by the arrow. The various topographic features are defined in the text.

# Hill Climbing Search

- The hill-climbing search algorithm (steepest-ascent version) is simply a loop that continually moves in the direction of increasing value. It terminates when it reaches a “peak”.
- Does not maintain a search tree, so the data structure for the current node need only record the state and the value of the objective function.

**function** HILL-CLIMBING(*problem*) **returns** a state that is a local maximum

*current*  $\leftarrow$  MAKE-NODE(*problem*.INITIAL-STATE)

**loop do**

*neighbor*  $\leftarrow$  a highest-valued successor of *current*

**if** *neighbor*.VALUE  $\leq$  *current*.VALUE **then return** *current*.STATE

*current*  $\leftarrow$  *neighbor*


**Figure 4.2** The hill-climbing search algorithm, which is the most basic local search technique. At each step the current node is replaced by the best neighbor; in this version, that means the neighbor with the highest VALUE, but if a heuristic cost estimate  $h$  is used, we would find the neighbor with the lowest  $h$ .

# Hill Climbing Search

- 8-queens problem
- The successors of a state are all possible states generated by moving a single queen to another square in the same column. The heuristic cost function  $h$  is the number of pairs of queens that are attacking each other.
- Hill-climbing algorithms typically choose randomly among the set of best successors if there is more than one.

18	12	14	13	13	12	14	14
14	16	13	15	12	14	12	16
14	12	18	13	15	12	14	14
15	14	14	♙	13	16	13	16
♙	14	17	15	♙	14	16	16
17	♙	16	18	15	♙	15	♙
18	14	♙	15	15	14	♙	16
14	14	13	17	12	14	12	18

(a)




(b)

**Figure 4.3** (a) An 8-queens state with heuristic cost estimate  $h = 17$ , showing the value of  $h$  for each possible successor obtained by moving a queen within its column. The best moves are marked. (b) A local minimum in the 8-queens state space; the state has  $h = 1$  but every successor has a higher cost.

# Hill Climbing Search

- Hill climbing is sometimes called **greedy local search** because it grabs a good neighbor state without thinking ahead about where to go next.
  - It turns out that greedy algorithms often perform quite well
- Hill climbing often gets stuck for the following reasons:
  - Local maxima
  - Ridges (屋脊)
  - Plateaux (can be flat local maximum or a **shoulder**)



**Figure 4.4** Illustration of why ridges cause difficulties for hill climbing. The grid of states (dark circles) is superimposed on a ridge rising from left to right, creating a sequence of local maxima that are not directly connected to each other. From each local maximum, all the available actions point downhill.

# Hill Climbing Search

- For the 8-queens problem, steepest-ascent hill climbing gets stuck 86% of the time, solving only 14% of problem instances. It works quickly, taking just 4 steps on average when it succeeds and 3 when it gets stuck—not bad for a state space with  $8^8 \approx 17$  million states.
- Might it not be a good idea to keep going—to allow a sideways move in the hope that the plateau is really a shoulder? The answer is usually yes. This raises the percentage of problem instances solved by hill climbing from 14% to 94%. Success comes at a cost: the algorithm averages roughly 21 steps for each successful instance and 64 for each failure.



# Hill Climbing Search

- **Stochastic hill climbing** chooses at random from among the uphill moves; the probability of selection can vary with the steepness of the uphill move.
- **First-choice hill climbing** implements stochastic hill climbing by generating successors randomly until one is generated that is better than the current state.
- **Random-restart hill climbing** conducts a series of hill-climbing searches from randomly generated initial states until a goal is found.
- The success of hill climbing depends very much on the shape of the state-space landscape.

# Simulated Annealing

- **Annealing** is the process used to temper or harden metals and glass by heating them to a high temperature and then gradually cooling them, thus allowing the material to reach a low-energy crystalline state.

```
function SIMULATED-ANNEALING(problem, schedule) returns a solution state
  inputs: problem, a problem
           schedule, a mapping from time to “temperature”

  current  $\leftarrow$  MAKE-NODE(problem.INITIAL-STATE)
  for  $t = 1$  to  $\infty$  do
     $T \leftarrow \text{schedule}(t)$ 
    if  $T = 0$  then return current
    next  $\leftarrow$  a randomly selected successor of current
     $\Delta E \leftarrow \text{next.VALUE} - \text{current.VALUE}$ 
    if  $\Delta E > 0$  then current  $\leftarrow$  next
    else current  $\leftarrow$  next only with probability  $e^{\Delta E/T}$ 
```

**Figure 4.5** The simulated annealing algorithm, a version of stochastic hill climbing where some downhill moves are allowed. Downhill moves are accepted readily early in the annealing schedule and then less often as time goes on. The *schedule* input determines the value of the temperature  $T$  as a function of time.

# Simulated Annealing

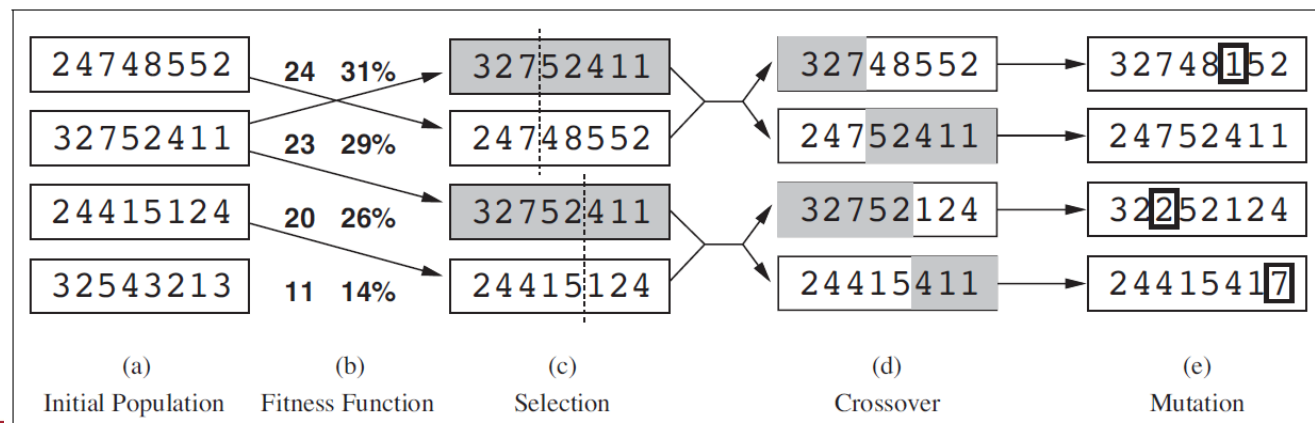
- Instead of picking the best move, however, it picks a random move. If the move improves the situation, it is always accepted. Otherwise, the algorithm accepts the move with some probability less than 1. The probability decreases exponentially with the “badness” of the move.
- The probability also decreases as the “temperature”  $T$  goes down: “bad” moves are more likely to be allowed at the start when  $T$  is high, and they become more unlikely as  $T$  decreases.

# Local Beam Search

- The **local beam search** algorithm keeps track of  $k$  states rather than just one.
- It begins with  $k$  randomly generated states. At each step, all the successors of all  $k$  states are generated. If any one is a goal, the algorithm halts. Otherwise, it selects the  $k$  best successors from the complete list and repeats.
- A local beam search seem to be nothing more than running  $k$  random restarts in parallel instead of in sequence. In a random-restart search, each search process runs independently of the others. In a local beam search, useful information is passed among the parallel search threads.
- **Stochastic beam search** chooses  $k$  successors at random, with the probability of choosing a given successor being an increasing function of its value.

# Genetic Algorithms


- A genetic algorithm (or GA) is a variant of stochastic beam search in which successor states are generated by combining two parent states rather than by modifying a single state.
- GAs begin with a set of  $k$  randomly generated states, called the **population**. Each state, or **individual**, is represented as a string over a finite alphabet—most commonly, a string of 0s and 1s.



**Figure 4.6** The genetic algorithm, illustrated for digit strings representing 8-queens states. The initial population in (a) is ranked by the fitness function in (b), resulting in pairs for mating in (c). They produce offspring in (d), which are subject to mutation in (e).

# Genetic Algorithms

- A **fitness** function should return higher values for better states. The probability of being chosen for reproducing is directly proportional to the fitness score.
- For each pair to be mated, a **crossover** point is chosen randomly from the positions in the string.



**Figure 4.7** The 8-queens states corresponding to the first two parents in Figure 4.6(c) and the first offspring in Figure 4.6(d). The shaded columns are lost in the crossover step and the unshaded columns are retained.

# Genetic Algorithms

- Finally, each location is subject to random mutation with a small independent probability. In the 8-queens problem, this corresponds to choosing a queen at random and moving it to a random square in its column.

```
function GENETIC-ALGORITHM(population, FITNESS-FN) returns an individual
  inputs: population, a set of individuals
           FITNESS-FN, a function that measures the fitness of an individual

  repeat
    new_population  $\leftarrow$  empty set
    for  $i = 1$  to SIZE(population) do
       $x \leftarrow$  RANDOM-SELECTION(population, FITNESS-FN)
       $y \leftarrow$  RANDOM-SELECTION(population, FITNESS-FN)
       $child \leftarrow$  REPRODUCE( $x, y$ )
      if (small random probability) then  $child \leftarrow$  MUTATE( $child$ )
      add  $child$  to new_population
    population  $\leftarrow$  new_population
  until some individual is fit enough, or enough time has elapsed
  return the best individual in population, according to FITNESS-FN
```

---

```
function REPRODUCE( $x, y$ ) returns an individual
  inputs:  $x, y$ , parent individuals

   $n \leftarrow$  LENGTH( $x$ );  $c \leftarrow$  random number from 1 to  $n$ 
  return APPEND(SUBSTRING( $x, 1, c$ ), SUBSTRING( $y, c + 1, n$ ))
```

**Figure 4.8** A genetic algorithm. The algorithm is the same as the one diagrammed in Figure 4.6, with one variation: in this more popular version, each mating of two parents produces only one offspring, not two.

# Local Search in Continuous Spaces

- Suppose we want to place three new airports anywhere in Romania, such that the sum of squared distances from each city on the map to its nearest airport is minimized.
- The state space is then defined by the coordinates of the airports:  $(x_1, y_1)$ ,  $(x_2, y_2)$ , and  $(x_3, y_3)$ . This is a six-dimensional space; we also say that states are defined by six **variables**.



# Local Search in Continuous Spaces

- Let  $C_i$  be the set of cities whose closest airport (in the current state) is airport  $i$ .

The objective function is

$$f(x_1, y_1, x_2, y_2, x_3, y_3) = \sum_{i=1}^3 \sum_{c \in C_i} (x_i - x_c)^2 + (y_i - y_c)^2$$

- To avoid continuous problems: **discretize** the neighborhood of each state. We can move only one airport at a time in either the  $x$  or  $y$  direction by a fixed amount  $\pm\delta$ . With 6 variables, this gives 12 possible successors for each state. We can then apply any of the local search algorithms described previously.

# Local Search in Continuous Spaces

- Many methods attempt to use the gradient of the landscape to find a maximum.

The gradient of the objective function is a vector  $\nabla f$  that gives the magnitude and direction of the steepest slope.

$$\nabla f = \left( \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial y_1}, \frac{\partial f}{\partial x_2}, \frac{\partial f}{\partial y_2}, \frac{\partial f}{\partial x_3}, \frac{\partial f}{\partial y_3} \right)$$

- In some cases, we can find a maximum by solving the equation  $\nabla f = 0$ . In many cases, however, this equation cannot be solved in closed form.

# Local Search in Continuous Spaces

- For example, with three airports, the expression for the gradient depends on what cities are closest to each airport in the current state. This means we can compute the gradient locally; for example,

$$\frac{\partial f}{\partial x_1} = 2 \sum_{c \in C_1} (x_i - x_c)$$

- Given a locally correct expression for the gradient, we can perform steepest-ascent hill climbing by updating the current state according to the formula

$$\mathbf{x} \leftarrow \mathbf{x} + \alpha \nabla f(\mathbf{x})$$

where  $\alpha$  is a small constant often called the **step size**.


# Local Search in Continuous Spaces

- If  $\alpha$  is too small, too many steps are needed; if  $\alpha$  is too large, the search could overshoot the maximum. The technique of **line search** tries to overcome this dilemma by extending the current gradient direction until  $f$  starts to decrease again.
- For many problems, the most effective algorithm is the **Newton–Raphson** method. This is a general technique for finding roots of functions—that is, solving equations of the form  $g(x)=0$ . It works by computing a new estimate for the root  $x$  according to Newton’s formula

$$x \leftarrow x - g(x)/g'(x)$$

# Introduction

- The gradient of  $f$  at  $\mathbf{x}_0$ , denoted by  $\nabla f(\mathbf{x}_0)$ , is orthogonal to the tangent vector to an arbitrary smooth curve passing through  $\mathbf{x}_0$  on the level set  $f(\mathbf{x}) = c$
- The direction of maximum rate of increase of a real-valued differentiable function at a point is orthogonal to the level set of the function through that point.
- The gradient acts in such a direction that for a given small displacement, the function  $f$  increases more in the direction of the gradient than in any other direction




# Newton's Method

- Newton's method for solving equations of the form  $g(x) = 0$  is also referred to as *Newton's method of tangents*.
- If we draw a tangent to  $g(x)$  at the given point  $x^{(k)}$ , then the tangent line intersects the  $x$ -axis at the point  $x^{(k+1)}$ , which we expect to be closer to the root  $x^*$  of  $g(x) = 0$ .
- Note that the slope of  $g(x)$  at

$$g'(x^{(k)}) = \frac{g(x^{(k)})}{x^{(k)} - x^{(k+1)}}$$

➔ 
$$x^{(k+1)} = x^{(k)} - \frac{g(x^{(k)})}{g'(x^{(k)})}$$



# Local Search in Continuous Spaces

- To find a maximum or minimum of  $f$ , we need to find  $x$  such that the gradient is zero (i.e.,  $\nabla f(x) = 0$ ). Thus,  $g(x)$  in Newton's formula becomes  $\nabla f(x)$ , and the update equation can be written in matrix–vector form as

$$\mathbf{x} \leftarrow \mathbf{x} - \mathbf{H}_f^{-1}(\mathbf{x}) \nabla f(\mathbf{x})$$

where  $\mathbf{H}_f(\mathbf{x})$  is the **Hessian** matrix of second derivatives, whose elements  $H_{ij}$  are given by  $\partial^2 f / \partial x_i \partial x_j$ .

- Local search methods suffer from local maxima, ridges, and plateaux in continuous state spaces just as much as in discrete spaces. Random restarts and simulated annealing can be used and are often helpful.

# Local Search in Continuous Spaces

- A **constrained optimization** problem is constrained if solutions must satisfy some hard constraints on the values of the variables.
- The best-known category is that of **linear programming** problems, in which constraints must be linear inequalities forming a **convex set** and the objective function is also linear.
- Linear programming is probably the most widely studied and broadly useful class of optimization problems. It is a special case of the more general problem of **convex optimization**, which allows the constraint region to be any convex region and the objective to be any function that is convex within the constraint region.



# Simple Examples of Linear Programs

- Formally, a linear program is an optimization problem of the form

$$\begin{array}{ll}\text{minimize} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{Ax} = \mathbf{b} \quad \mathbf{x} \geq \mathbf{0}\end{array}$$

where  $\mathbf{c} \in R^n$ ,  $\mathbf{b} \in R^m$ ,  $\mathbf{A} \in R^{m \times n}$ . The vector inequality  $\mathbf{x} \geq \mathbf{0}$  means that each component of  $\mathbf{x}$  is nonnegative.

- Several variations of this problem are possible. For example, we can maximize, or the constraints may be in the form of inequalities, such as  $\mathbf{Ax} \geq \mathbf{b}$  or  $\mathbf{Ax} \leq \mathbf{b}$ . In fact, these variations can all be rewritten into the standard form shown above.

# Example

- A manufacturer produces four different products:  $X_1, X_2, X_3, X_4$  there are three inputs to this production process: labor in person-weeks, kilograms of raw material A, and boxes of raw material B. Each product has different input requirements. In determining each week's production schedule, the manufacturer cannot use more than the available amounts of labor and the two raw materials. The relevant information is presented in this table. Every production decision must satisfy the restrictions on the availability of inputs. These constraints can be written using the data in this table.

$$x_1 + 2x_2 + x_3 + 2x_4 \leq 20$$

$$6x_1 + 5x_2 + 3x_3 + 2x_4 \leq 100$$


$$3x_1 + 4x_2 + 9x_3 + 12x_4 \leq 75$$

Inputs	Products				Input Availabilities
	$X_1$	$X_2$	$X_3$	$X_4$	
man weeks	1	2	1	2	20
kilograms of material A	6	5	3	2	100
boxes of material B	3	4	9	12	75
production levels	$x_1$	$x_2$	$x_3$	$x_4$	


# Searching with Partial Observations

- A robot is placed in the maze-like environment. It is equipped with four sonar sensors that tell whether there is an obstacle in each of the four compass directions.
- Assume that the sensors give perfectly correct data, and the robot has a correct map of the environment. But unfortunately the robot's navigational system is broken, so when it executes a *Move* action, it moves randomly to one of the adjacent squares. The robot's task is to determine its current **location**.

# Searching with Partial Observations



(a) Possible locations of robot after  $E_1 = \text{NSW}$



(b) Possible locations of robot After  $E_1 = \text{NSW}, E_2 = \text{NS}$

**Figure 4.18** Possible positions of the robot,  $\odot$ , (a) after one observation  $E_1 = \text{NSW}$  and (b) after a second observation  $E_2 = \text{NS}$ . When sensors are noiseless and the transition model is accurate, there are no other possible locations for the robot consistent with this sequence of two observations.

# Online Searching Agents with Unknown Environments


- So far we have concentrated on agents that use **offline search** algorithms. They compute a complete solution before setting foot in the real world and then execute the solution.
- In contrast, an **online search** agent interleaves computation and action: first it takes an action, then it observes the environment and computes the next action.
- The canonical example of online search is a robot that is placed in a new building and must explore it to build a map that it can use for getting from  $A$  to  $B$ .

# Online Searching Agents with Unknown Environments

- Online search algorithms
  - We stipulate (規定) that the agent knows only the following
    - $\text{ACTIONS}(s)$ , which returns a list of actions allowed in state  $s$ ;
    - The step-cost function  $c(s, a, s')$ —note that this cannot be used until the agent knows that  $s'$  is the outcome; and
    - $\text{GOAL-TEST}(s)$ .
  - The agent cannot determine  $\text{RESULT}(s, a)$  except by actually being in  $s$  and doing  $a$ .

# Online Searching Agents with Unknown Environments

- Online search algorithms
  - In the maze problem shown in Figure 4.19, the agent does not know that going *Up* from (1,1) leads to (1,2); nor, having done that, does it know that going *Down* will take it back to (1,1).



# Online Searching Agents with Unknown Environments

- Online search algorithms
  - Finally, the agent might have access to an admissible heuristic function  $h(s)$  that estimates the distance from the current state to a goal state. For example, in Figure 4.19, the agent might know the location of the goal and be able to use the Manhattan-distance heuristic.



# Online Searching Agents with Unknown Environments

- Online search algorithms
  - Typically, the agent's objective is to reach a goal state while minimizing cost. The cost is the total path cost of the path that the agent actually travels. It is common to compare this cost with the path cost of the path the agent would follow if it knew the search space in advance. This is called the **competitive ratio**; we would like it to be as small as possible.

# Online Searching Agents with Unknown Environments

- Online search agents
  - After each action, an online agent receives a percept telling it what state it has reached; from this info., it can augment its map of the environment.
  - The current map is used to decide where to go next. This interleaving of planning and action means that online search algorithms are quite different from the offline search algorithms we have seen previously.
  - To avoid traveling all the way across the tree to expand the next node, an online algorithm better expands nodes in a local order. DFS has exactly this property.

# Online Searching Agents with Unknown Environments

- Online local search
  - Like depth-first search, hill-climbing search has the property of locality in its node expansions. In fact, because it keeps just one current state in memory, hill-climbing search is already an online search algorithm!  
Unfortunately, it is not very useful in its simplest form because it leaves the agent sitting at local maxima with nowhere to go.