

G5BAIM **Artificial Intelligence Methods**

Dr. Rong Qu
Simulated Annealing



Simulated Annealing

- Motivated by the physical annealing process
- Material is heated and slowly cooled into a uniform structure
- Simulated annealing mimics this process
- The first SA algorithm was developed in 1953 (Metropolis)

Simulated Annealing

- **Kirkpatrick (1982) applied SA to optimisation problems**
 - Kirkpatrick, S , Gelatt, C.D., Vecchi, M.P. 1983. *Optimization by Simulated Annealing*. Science, vol 220, No. 4598, pp 671-680

The Problem with Hill Climbing

- Gets stuck at local minima
- Possible solutions
 - Try several runs, starting at different positions
 - Increase the size of the neighborhood (e.g. in TSP try 3-opt rather than 2-opt)

Simulated Annealing

- Compared to hill climbing the main difference is that SA allows downwards steps
- Simulated annealing also differs from hill climbing in that a move is selected at random and then decides whether to accept it
- In SA better moves are always accepted. Worse moves are not.

To accept or not to accept?

- The law of thermodynamics states that at temperature, t , the probability of an increase in energy of magnitude, δE , is given by

$$P(\delta E) = \exp(-\delta E / kt)$$

- Where k is a constant known as Boltzmann's constant

To accept or not to accept - SA?

$$P = \exp(-c/t) > r$$

Where

- n c is change in the evaluation function
- n t the current temperature
- n r is a random number between 0 and 1

Example

To accept or not to accept - SA?

Change	Temp	exp(-C/T)	Change	Temp	exp(-C/T)
0.2	0.95	0.810157735	0.2	0.1	0.135335283
0.4	0.95	0.656355555	0.4	0.1	0.018315639
0.6	0.95	0.53175153	0.6	0.1	0.002478752
0.8	0.95	0.430802615	0.8	0.1	0.000335463

* Need to use a scientific calculator to calculate exp()

To accept or not to accept - SA?

- n The probability of accepting a worse state is a function of both the temperature of the system and the change in the cost function
- n As the temperature decreases, the probability of accepting worse moves decreases
- n If $t=0$, no worse moves are accepted (i.e. hill climbing)

SA Algorithm

- The most common way of implementing an SA algorithm is to implement hill climbing with an accept function and modify it for SA
- The example shown here is taken from Russell/Norvig ([Artificial Intelligence: A Modern Approach](#))

SA Algorithm

- n **Function** SIMULATED-ANNEALING(*Problem*, *Schedule*) **returns** a solution state
- n **Inputs:** *Problem*, a problem
Schedule, a mapping from time to temperature
Local Variables: *Current*, a node
Next, a node
T, a "temperature" controlling the probability of downward steps
- n *Current* = MAKE-NODE(INITIAL-STATE[*Problem*])

SA Algorithm

```

For t = 1 to ∞ do
  T = Schedule[t]
  If T = 0 then return Current
  Next = a randomly selected successor of Current
  ΔE = VALUE[Next] - VALUE[Current]
  if ΔE > 0 then Current = Next
  else Current = Next only with probability exp(-ΔE/T)
  
```



SA Algorithm

- n The cooling schedule is *hidden* in this algorithm - but it is important (more later)
- n The algorithm assumes that annealing will continue until temperature is zero - this is not necessarily the case



SA Cooling Schedule

- n Starting Temperature
- n Final Temperature
- n Temperature Decrement
- n Iterations at each temperature



SA Cooling Schedule - Starting Temperature

- n Starting Temperature
 - n Must be *hot* enough to allow moves to *almost* neighbourhood state (else we are in danger of implementing hill climbing)
 - n Must *not* be so hot that we conduct a random search for a period of time
 - n Problem is finding a suitable starting temperature



SA Cooling Schedule - Starting Temperature

- n Starting Temperature - Choosing
 - n If we know the maximum change in the cost function we can use this to estimate
 - n Start high, reduce quickly until about 60% of worse moves are accepted. Use this as the starting temperature
 - n Heat rapidly until a certain percentage are accepted the start cooling



SA Cooling Schedule - Final Temperature

- n Final Temperature - Choosing
 - n It is usual to let the temperature decrease until it reaches zero
However, this can make the algorithm run for a lot longer, especially when a geometric cooling schedule is being used
 - n In practise, it is not necessary to let the temperature reach zero because the chances of accepting a worse move are almost the same as the temperature being equal to zero



SA Cooling Schedule - Final Temperature

- n Final Temperature - Choosing
 - n Therefore, the stopping criteria can either be a suitably low temperature or when the system is "frozen" at the current temperature (i.e. no better or worse moves are being accepted)

SA Cooling Schedule - Temperature Decrement

- n Temperature Decrement
 - n Theory states that we should allow enough iterations at each temperature so that the system stabilises at that temperature
 - n Unfortunately, theory also states that the number of iterations at each temperature to achieve this might be exponential to the problem size

SA Cooling Schedule - Temperature Decrement

- n Temperature Decrement
 - n We need to compromise
 - n We can either do this by doing a large number of iterations at a few temperatures, a small number of iterations at many temperatures or a balance between the two

SA Cooling Schedule - Temperature Decrement


- n Temperature Decrement
 - n Linear
 - $temp = temp - x$
 - n Geometric
 - $temp = temp * a$
 - Experience has shown that a should be between 0.8 and 0.99, with better results being found in the higher end of the range. Of course, the higher the value of a , the longer it will take to decrement the temperature to the stopping criterion

SA Cooling Schedule - Iterations

- n Iterations at each temperature
 - n A constant number of iterations at each temperature
 - n Another method, first suggested by (Lundy, 1986) is to only do one iteration at each temperature, but to decrease the temperature *very* slowly.

SA Cooling Schedule - Iterations

- n Iterations at each temperature
 - n The formula used by Lundy is
 - $t = t / (1 + \beta t)$
 - n where β is a suitably small value



SA Cooling Schedule - Iterations

- n Iterations at each temperature
 - n An alternative is to dynamically change the number of iterations as the algorithm progresses
 - At lower temperatures it is important that a large number of iterations are done so that the local optimum can be fully explored
 - At higher temperatures, the number of iterations can be less



Problem Specific Decisions

- n The cooling schedule is all about SA but there are other decisions which we need to make about the problem
- n These decisions are not just related to SA



Problem Specific Decisions - Cost Function

- n The evaluation function is calculated at every iteration
- n Often the cost function is the most expensive part of the algorithm



Problem Specific Decisions - Cost Function

- n Therefore
 - n We need to evaluate the cost function as efficiently as possible
 - n Use Delta Evaluation
 - n Use Partial Evaluation



Problem Specific Decisions - Cost Function

- n If possible, the cost function should also be designed so that it can lead the search
 - n One way of achieving this is to avoid cost functions where many states return the same value
This can be seen as representing a plateau in the search space which the search has no knowledge about which way it should proceed
 - n Bin Packing



Problem Specific Decisions - Cost Function Example

- n Bin Packing
 - n A number of items, a number of bins
 - n Objective
 - n As many items as possible
 - n As less bins as possible
 - n Other objectives depending on the problems



Problem Specific Decisions - Cost Function Example

- n Bin Packing
 - n Cost function?
 - n a) number of bins
 - n b) number of items
 - n c) both a) and b)
 - n How about there are weights for the items?

Problem Specific Decisions - Cost Function

- n Many cost functions cater for the fact that some solutions are illegal. This is typically achieved using constraints
 - n Hard Constraints : these constraints cannot be violated in a feasible solution
 - n Soft Constraints : these constraints should, ideally, not be violated but, if they are, the solution is still feasible
 - n Examples: bin packing, timetabling

Problem Specific Decisions - Cost Function

- n Hard constraints are given a large weighting. The solutions which violate those constraints have a high cost function
- n Soft constraints are weighted depending on their importance
- n Weightings can be dynamically changed as the algorithm progresses. This allows hard constraints to be accepted at the start of the algorithm but rejected later

Problem Specific Decisions - Neighbourhood

- n How do you move from one state to another?
- n When you are in a certain state, what other states are reachable?
 - n Examples: bin packing, timetabling

Problem Specific Decisions - Neighbourhood

- n Some results have shown that the neighbourhood structure should be symmetric. That is, if you move from state i to state j then it must be possible to move from state j to state i
- n However, a weaker condition can hold in order to ensure convergence.
- n Every state must be *reachable* from every other. Therefore, it is important, when thinking about your problem to ensure that this condition is met

Problem Specific Decisions – Search space

- n The smaller the search space, the easier the search will be
- n If we define cost function such that infeasible solutions are accepted, the search space will be increased
- n As well as keeping the search space small, also keep the neighbourhood small

Problem Specific Decisions

- n Search space - small
 - n large size of neighbourhood
 - n search is not restricted
- n Cost function - easy to calculate
 - n consider infeasible solutions
- n Overall aim
 - n Make the most use of each iteration, whilst trying to ensure good quality solution

Problem Specific Decisions - Performance

- n What is performance?
 - n Quality of the solution returned
 - n Time taken by the algorithm
- n We already have the problem of finding suitable SA parameters (cooling schedule)

Problem Specific Decisions - Performance

- n Improving Performance - Initialisation
 - n Start with a random solution and let the annealing process improve on that.
 - n Might be better to start with a solution that has been heuristically built (e.g. for the TSP problem, start with a greedy search)

Problem Specific Decisions - Performance

- n Improving Performance - Hybridisation
 - n or *memetic algorithms*
 - n Combine two search algorithms
 - n Relatively new research area

Problem Specific Decisions - Performance

- n Improving Performance - Hybridisation
 - n Often a population based search strategy is used as the primary search mechanism and a local search mechanism is applied to move each individual to a local optimum
 - n It may be possible to apply some heuristic to a solution in order to improve it

SA Modifications - Acceptance Probability

- n The probability of accepting a worse move is normally based on the physical analogy (based on the Boltzmann distribution)
- n But is there any reason why a different function will not perform better for all, or at least certain, problems?

SA Modifications - Acceptance Probability

- n Why should we use a different acceptance criteria?
 - n The one proposed does not work. Or we suspect we might be able to produce better solutions
 - n The exponential calculation is computationally expensive.
 - n (Johnson, 1991) found that the acceptance calculation took about one third of the computation time

SA Modifications - Acceptance Probability

- Johnson experimented with

$$P(\delta) = 1 - \delta/t$$
- This approximates the exponential

Iteration	Temperature	Energy	Acceptance Probability
1	1000	1000	0.000
2	900	900	0.011
3	800	800	0.022
4	700	700	0.033
5	600	600	0.044
6	500	500	0.056
7	400	400	0.067
8	300	300	0.078
9	200	200	0.089
10	100	100	0.100
11	100	100	0.100
12	100	100	0.100
13	100	100	0.100
14	100	100	0.100
15	100	100	0.100
16	100	100	0.100
17	100	100	0.100
18	100	100	0.100
19	100	100	0.100
20	100	100	0.100
21	100	100	0.100
22	100	100	0.100
23	100	100	0.100
24	100	100	0.100
25	100	100	0.100
26	100	100	0.100
27	100	100	0.100
28	100	100	0.100
29	100	100	0.100
30	100	100	0.100
31	100	100	0.100
32	100	100	0.100
33	100	100	0.100
34	100	100	0.100
35	100	100	0.100
36	100	100	0.100
37	100	100	0.100
38	100	100	0.100
39	100	100	0.100
40	100	100	0.100
41	100	100	0.100
42	100	100	0.100
43	100	100	0.100
44	100	100	0.100
45	100	100	0.100
46	100	100	0.100
47	100	100	0.100
48	100	100	0.100
49	100	100	0.100
50	100	100	0.100

SA Modifications - Acceptance Probability

- A better approach was found by building a look-up table of a set of values over the range δ/t
- During the course of the algorithm δ/t was rounded to the nearest integer and this value was used to access the look-up table
- This method was found to speed up the algorithm by about a third with no significant effect on solution quality

SA Modifications - Cooling

- If you plot a typical cooling schedule you are likely to find that at high temperatures many solutions are accepted
- If you start at too high a temperature a random search is emulated and until the temperature cools sufficiently any solution can be reached and could have been used as a starting position

SA Modifications - Cooling

- At lower temperatures, a plot of the cooling schedule, is likely to show that very few worse moves are accepted; almost making simulated annealing emulate hill climbing

SA Modifications - Cooling

- Taking this one stage further, we can say that simulated annealing does most of its work during the middle stages of the cooling schedule
- (Connolly, 1990) suggested annealing at a constant temperature

SA Modifications - Cooling

- But what temperature?
- It must be high enough to allow movement but not so low that the system is frozen
- But, the optimum temperature will vary from one type of problem to another and also from one instance of a problem to another instance of the same problem



SA Modifications - Cooling

- n One solution to this problem is to spend some time searching for the optimum temperature and then stay at that temperature for the remainder of the algorithm
- n The final temperature is chosen as the temperature that returns the best cost function during the search phase



SA Modifications - Neighbourhood

- n The neighbourhood of any move is normally the same throughout the algorithm but...
- n The neighbourhood could be changed as the algorithm progresses
 - n For example, a different neighbourhood can be used to help jumping from local optimal



SA Modifications - Cost Function

- n The cost function is calculated at every iteration of the algorithm
- n Various researchers (e.g. Burke, 1999) have shown that the cost function can be responsible for a large proportion of the execution time of the algorithm
- n Some techniques have been suggested which aim to alleviate this problem



SA Modifications - Cost Function

- n (Rana, 1996) - Coors Brewery
 - n GA but could be applied to SA
 - n The evaluation function is approximated (one tenth of a second)
 - n Potentially good solutions are fully evaluated (three minutes)



SA Modifications - Cost Function

- n (Ross, 1994) uses delta evaluation on the timetabling problem
 - n Instead of evaluating every timetable as only small changes are being made between one timetable and the next, it is possible to evaluate just the changes and update the previous cost function using the result of that calculation



SA Modifications - Cost Function

- n (Burke, 1999) uses a cache
 - n The cache stores cost functions (partial and complete) that have already been evaluated
 - n They can be retrieved from the cache rather than having to go through the evaluation function again

Summary

- n SA basics
 - n Acceptance criteria
 - n Cooling schedule
- n Problem specific decisions
 - n Cost function
 - n Neighborhood
 - n Performance (initialisation, hybridisation)
- n SA modifications

G5BAIM ***Artificial Intelligence Methods***

Rong Qu

End of Simulated Annealing

