

Local Search (Ch. 4-4.1)



Announcements

HW1 late penalty

Test next week: will cover HW 1 & 2

-Ch 1-3 (up to A* and admissibility)

-Open book/notes

Local search

Before we tried to find a path from the start state to a goal state using a “fringe” set

Now we will look at algorithms that do not care about a “fringe”, but just neighbors

Some problems, may not have a clear “best” goal, yet we have some way of evaluating the state (how “good” is a state)

Local search

Today we will talk about 4 (more) algorithms:

1. Hill climbing
2. Simulated annealing
3. Beam search
4. Genetic algorithms

All of these will only consider neighbors while looking for a goal

Local search

General properties of local searches:

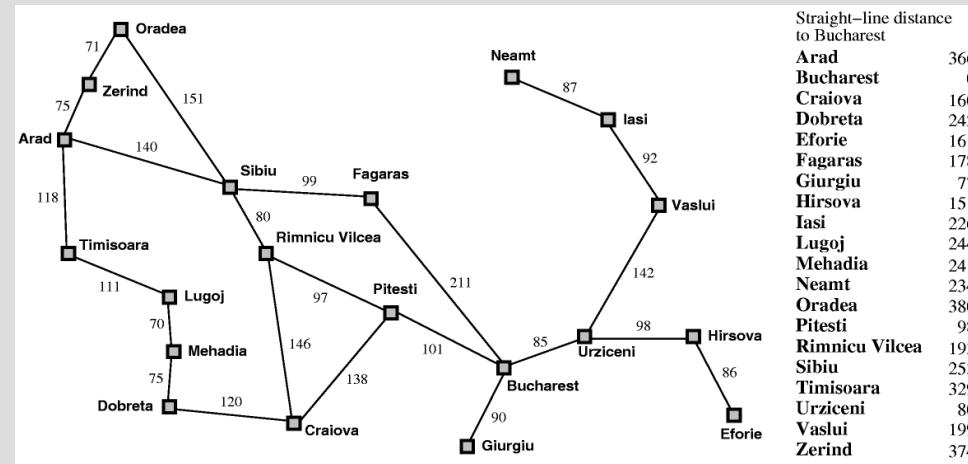
- Fast and low memory
- Can find “good” solutions if can estimate state value
- Hard to find “optimal” path

In general these types of searches are used if the tree is too big to find a real “optimal” solution

Hill climbing

Remember greedy best-first search?

1. Pick from neighbor with best heuristic
2. Repeat 1...



Hill climbing is only a slight variation:

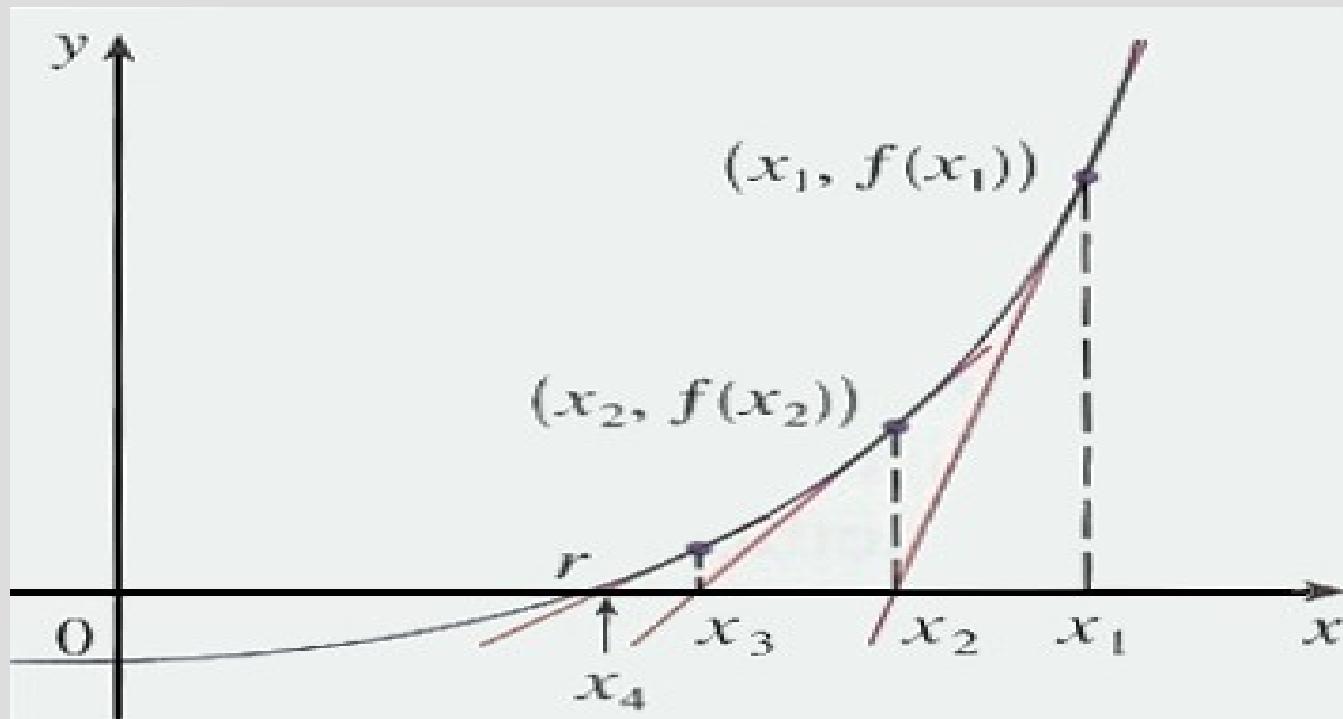
1. Pick best between: yourself and child
2. Repeat 1...

What are the pros and cons of this?

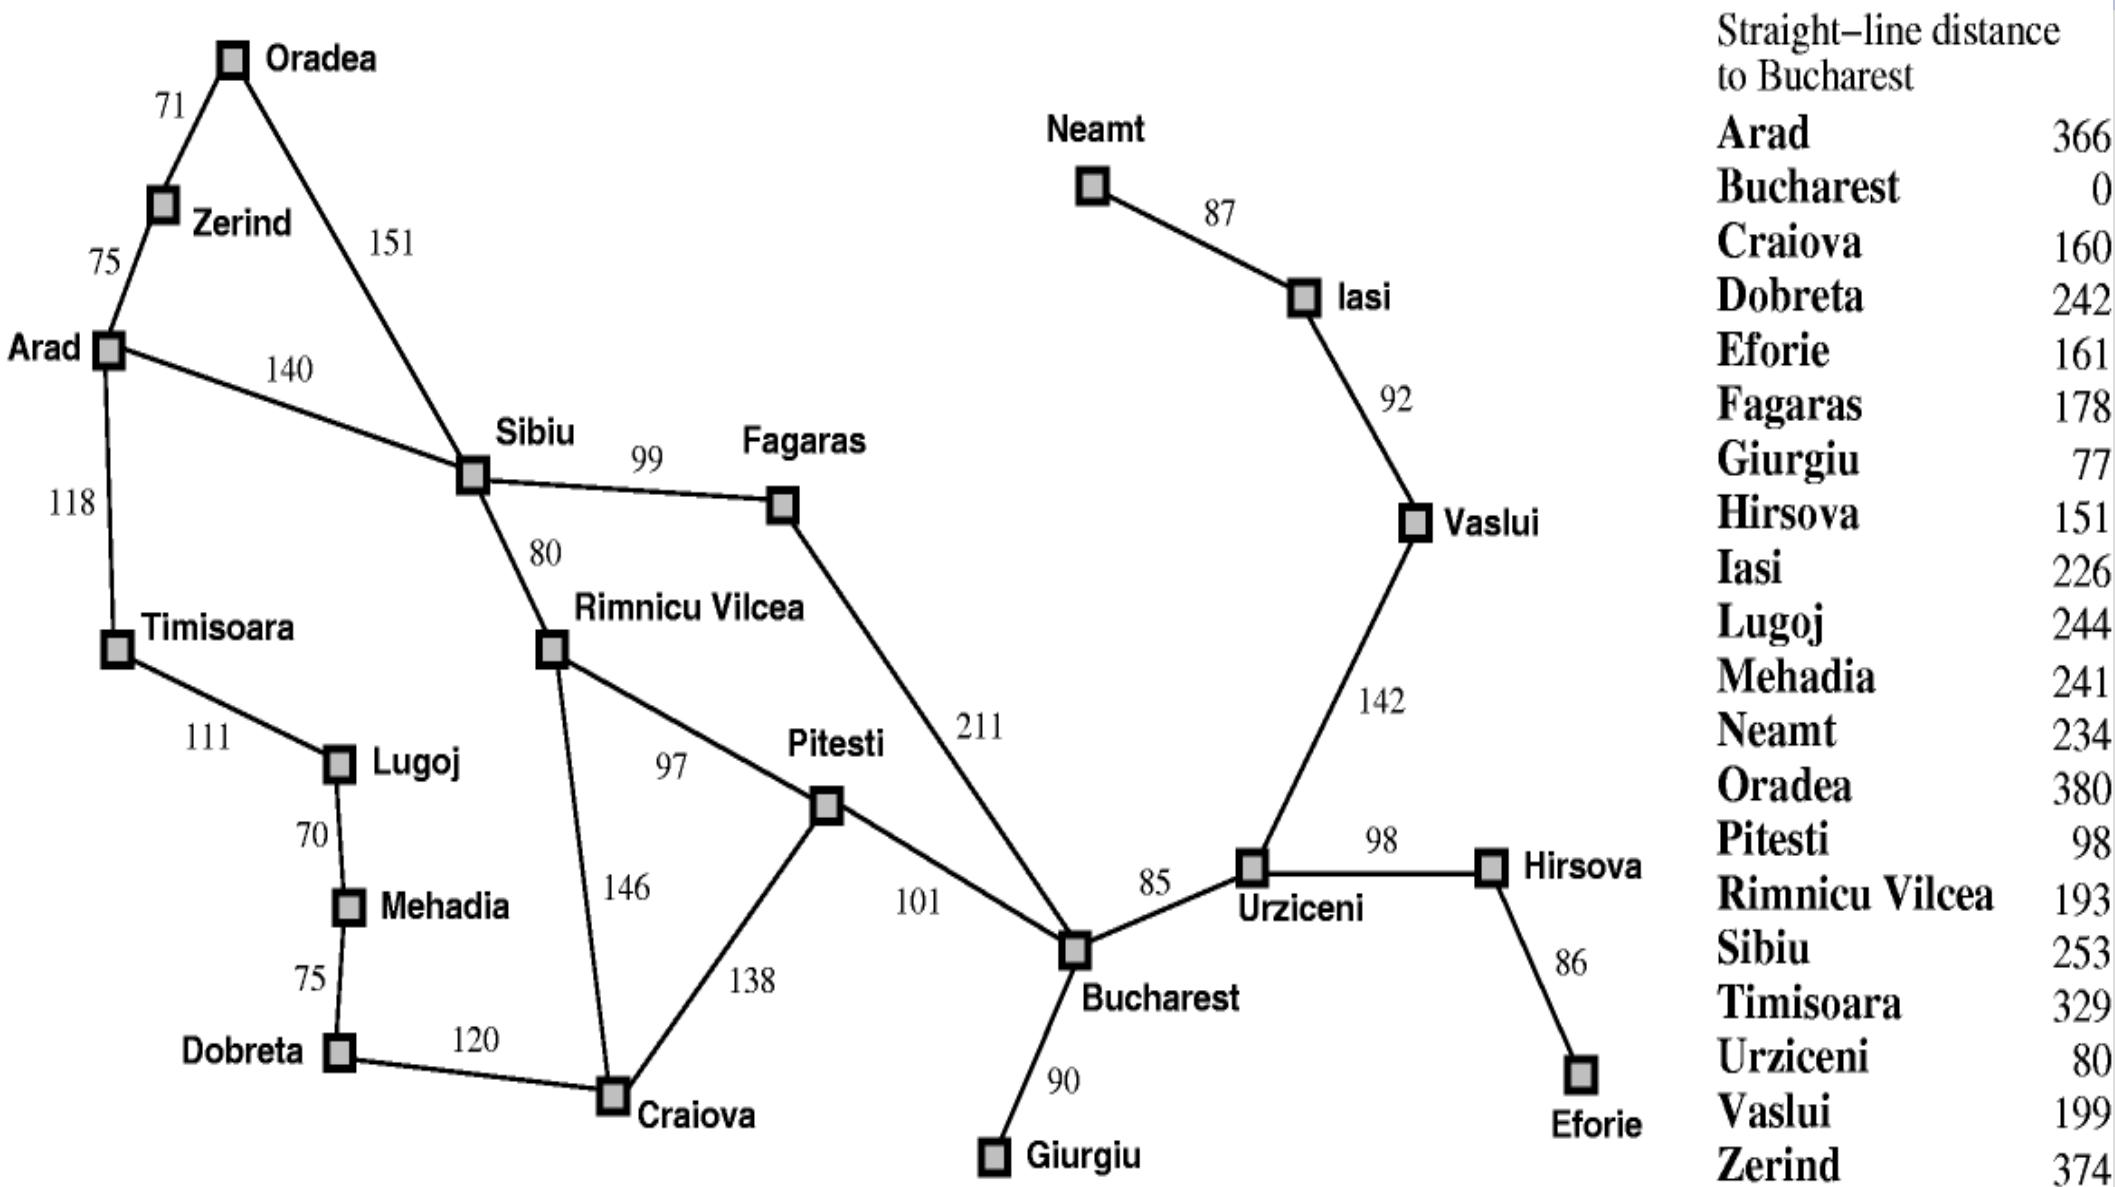
Hill climbing

This actually works surprisingly well, if getting “close” to the goal is sufficient (and actions are not too restrictive)

Newton's method:



Hill climbing



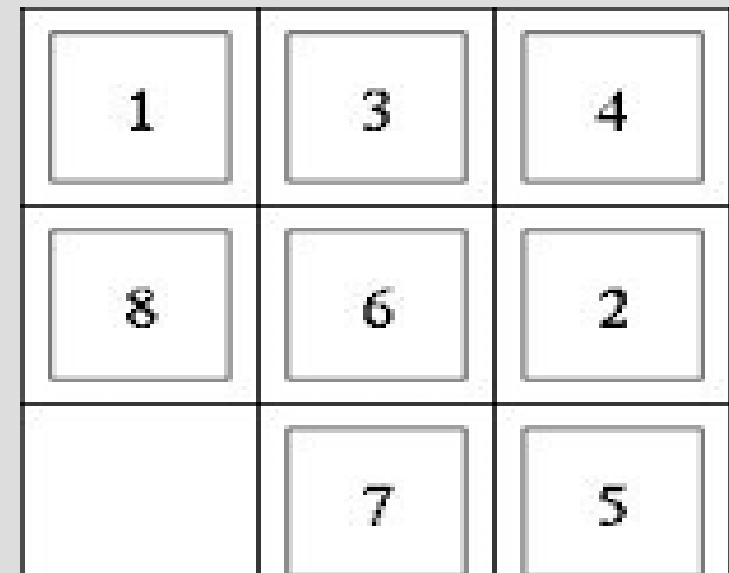
Hill climbing

For the 8-puzzles we had 2 (consistent) heuristics:

h_1 - number of mismatched pieces

h_2 - \sum Manhattan distance from number's current to goal position

Let's try hill climbing this problem!



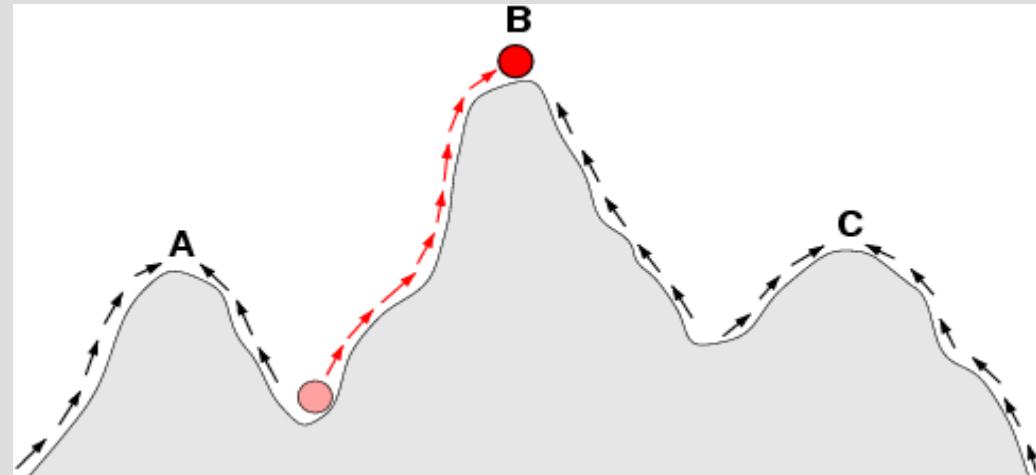
Hill climbing

Can get stuck in:

- Local maximum
- Plateau/shoulder

Local maximum will have a range of attraction around it

Can get an infinite loop in a plateau if not careful (step count)



Hill climbing

To avoid these pitfalls, most local searches incorporate some form of randomness

Hill search variants:

Stochastic hill climbing - choose random move from better solutions

Random-restart hill search - run hill search until maximum found (or looping), then start at another random spot and repeat

Simulated annealing

The idea behind simulated annealing is we act more random at the start (to “explore”), then take greedy choices later

<https://www.youtube.com/watch?v=qfD3cmQbn28>

An analogy might be a hard boiled egg:

1. To crack the shell you hit rather hard (not too hard!)
2. You then hit lightly to create a cracked area around first
3. Carefully peal the rest



Simulated annealing

The process is:

1. Pick random action and evaluation result
2. If result better than current, take it
3. If result worse accept probabilistically
4. Decrease acceptance chance in step 3
5. Repeat...

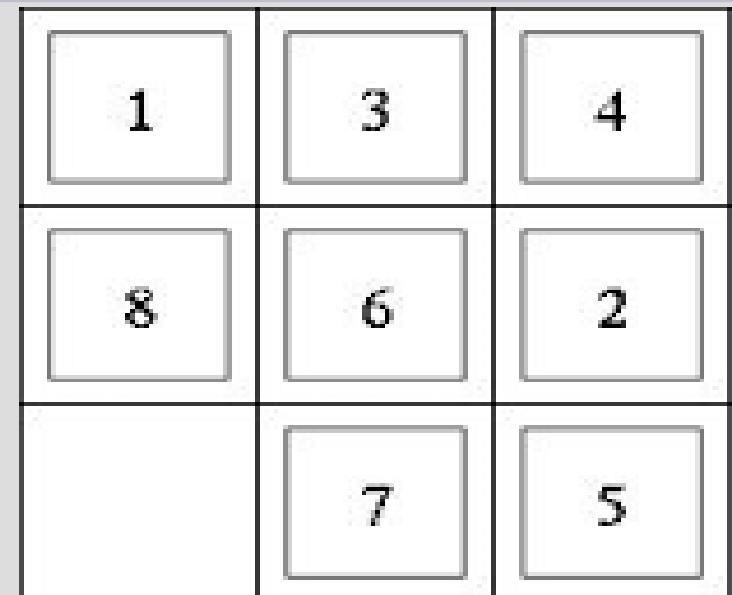
(see: SAacceptance.cpp)

Specifically, we track some “temperature” T :

3. Accept with probability: $e^{\frac{result - current}{T}}$
4. Decrease T (linear? hard to find best...)

Simulated annealing

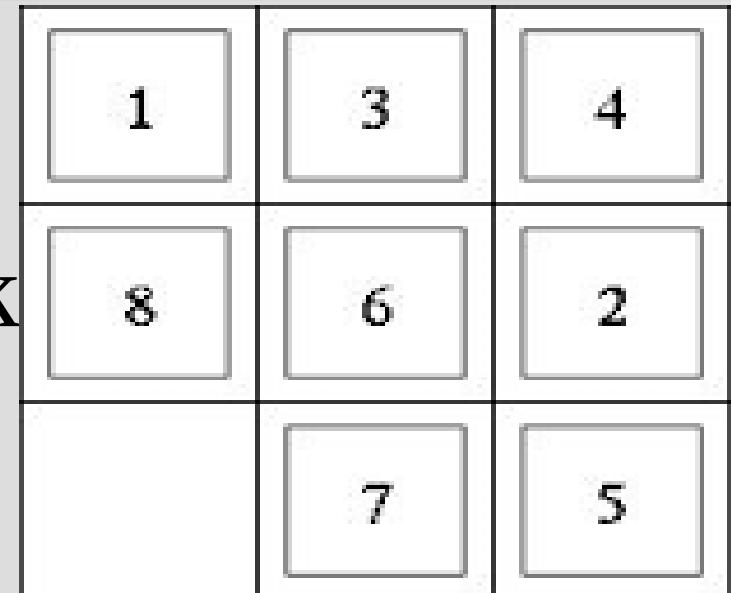
Let's try SA on 8-puzzle:



Simulated annealing

Let's try SA on 8-puzzle:

This example did not work well, but probably due to the temperature handling



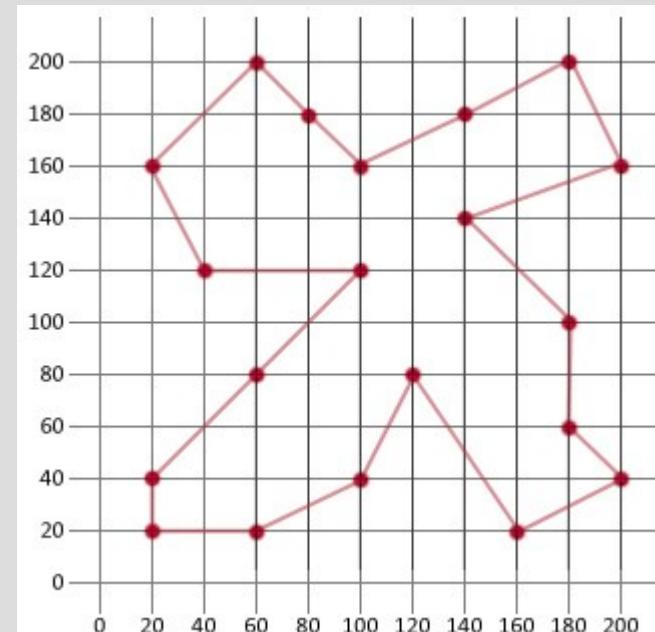
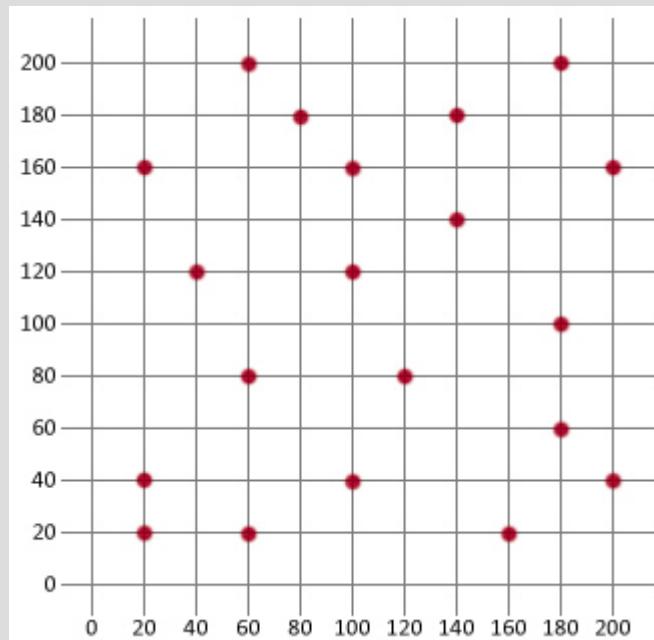
We want the temperature to be fairly high at the start (to move around the graph)

The hard part is slowly decreasing it over time

Simulated annealing

SA does work well on the traveling salesperson problem

(see: tsp.zip)



Local beam search

Beam search is similar to hill climbing, except we track multiple states simultaneously

Initialize: start with K random nodes

1. Find all children of the K nodes
2. Add children and K nodes to pool, pick best
3. Repeat...

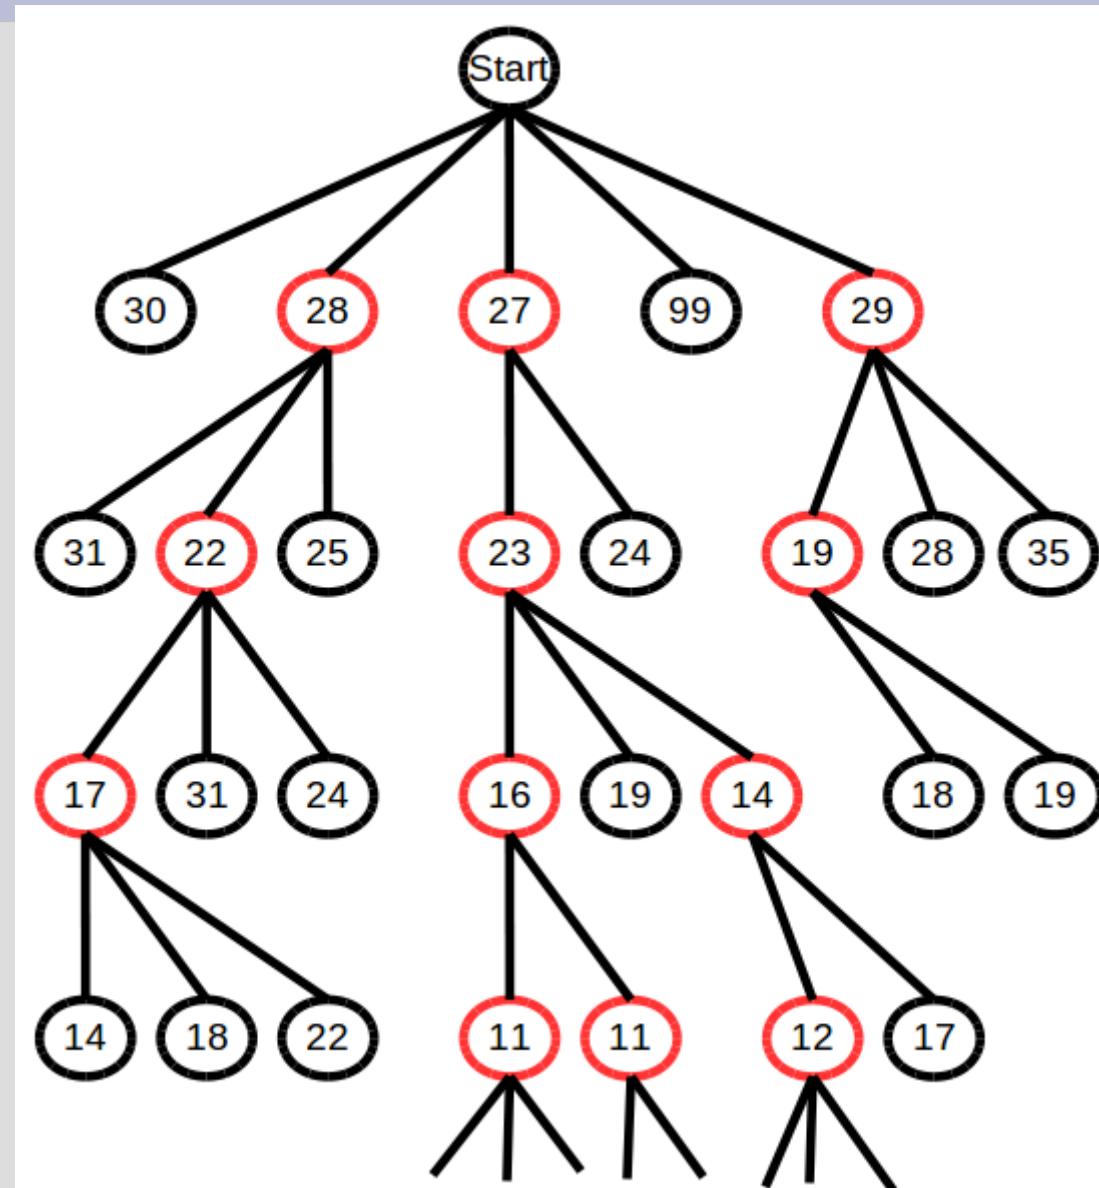
Unlike previous approaches, this uses more memory to better search “hopeful” options

Local beam search

Beam search with
3 beams

Pick best 3 options
at each depth to
expand

Stop like hill-climb
(next pick is
same as last pick)



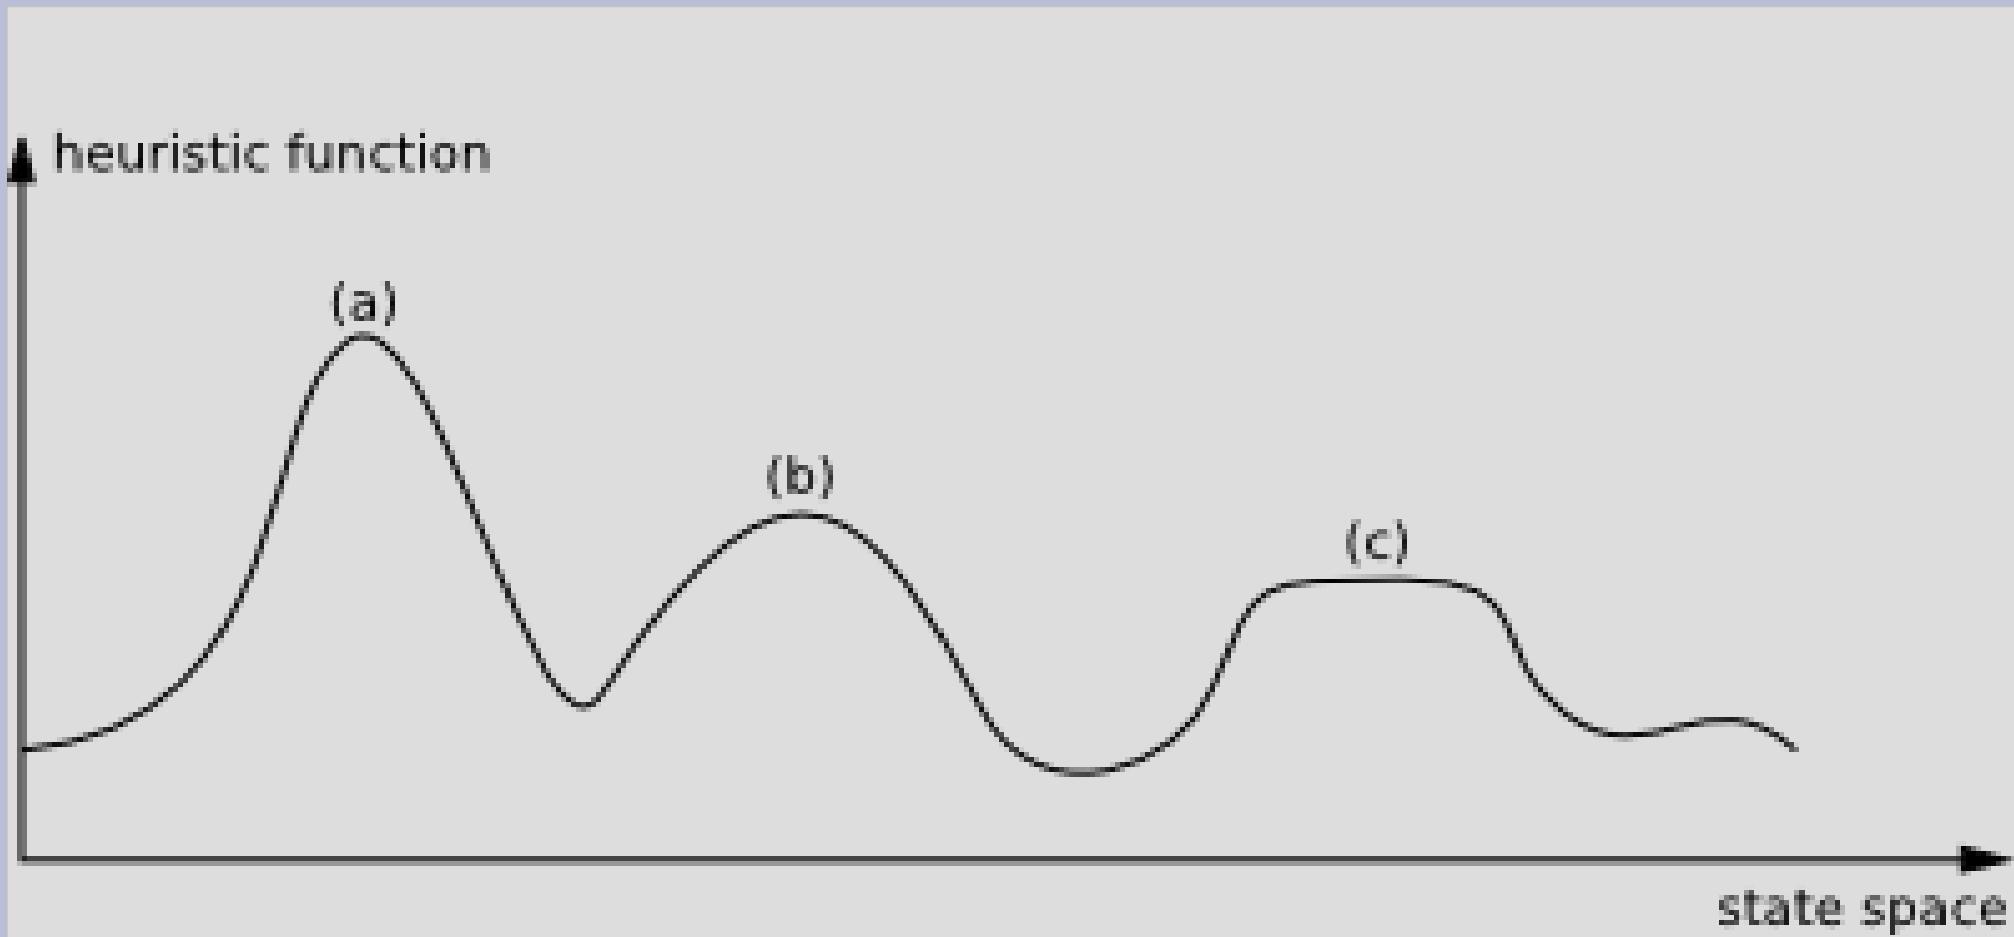
Local beam search

However, the basic version of beam search can get stuck in local maximum as well

To help avoid this, stochastic beam search picks children with probability relative to their values

This is different than hill climbing with K restarts as better options get more consideration than worse ones

Local beam search



Genetic algorithms

Nice examples of GAs:

http://rednuht.org/genetic_cars_2/
<http://boxcar2d.com/>

BoxCar 2D

Home | Designer | Best Cars | Forum | News | FAQ | The Algorithm | Versions | Contact
Computation Intelligence Car Evolution Using Box2D Physics (v3.2)

60 fps average
Physics step: 1 ms (833 fps)
18 MB used

Input Seed / Choose Terrain 139
Generation: 4 Max Score: 139.5

69

139.5

84.4

3.8

0.3

3.4

1.8

10

3.2

22.2

11.6

0:24

0:10

0:01

0:02

0:00

0:05

0:02

0:10

0:03

Watch Leader

Car	Score	Time
0	139.5	0:24
1	84.4	0:10
2	3.8	0:01
3	0.3	0:02
4	3.4	0:02
5	1.8	0:00
6	10	0:05
7	3.2	0:02
8	22.2	0:10
9	11.6	0:03

Up
Next
Down
Copy Current
Copy Best

View top replay

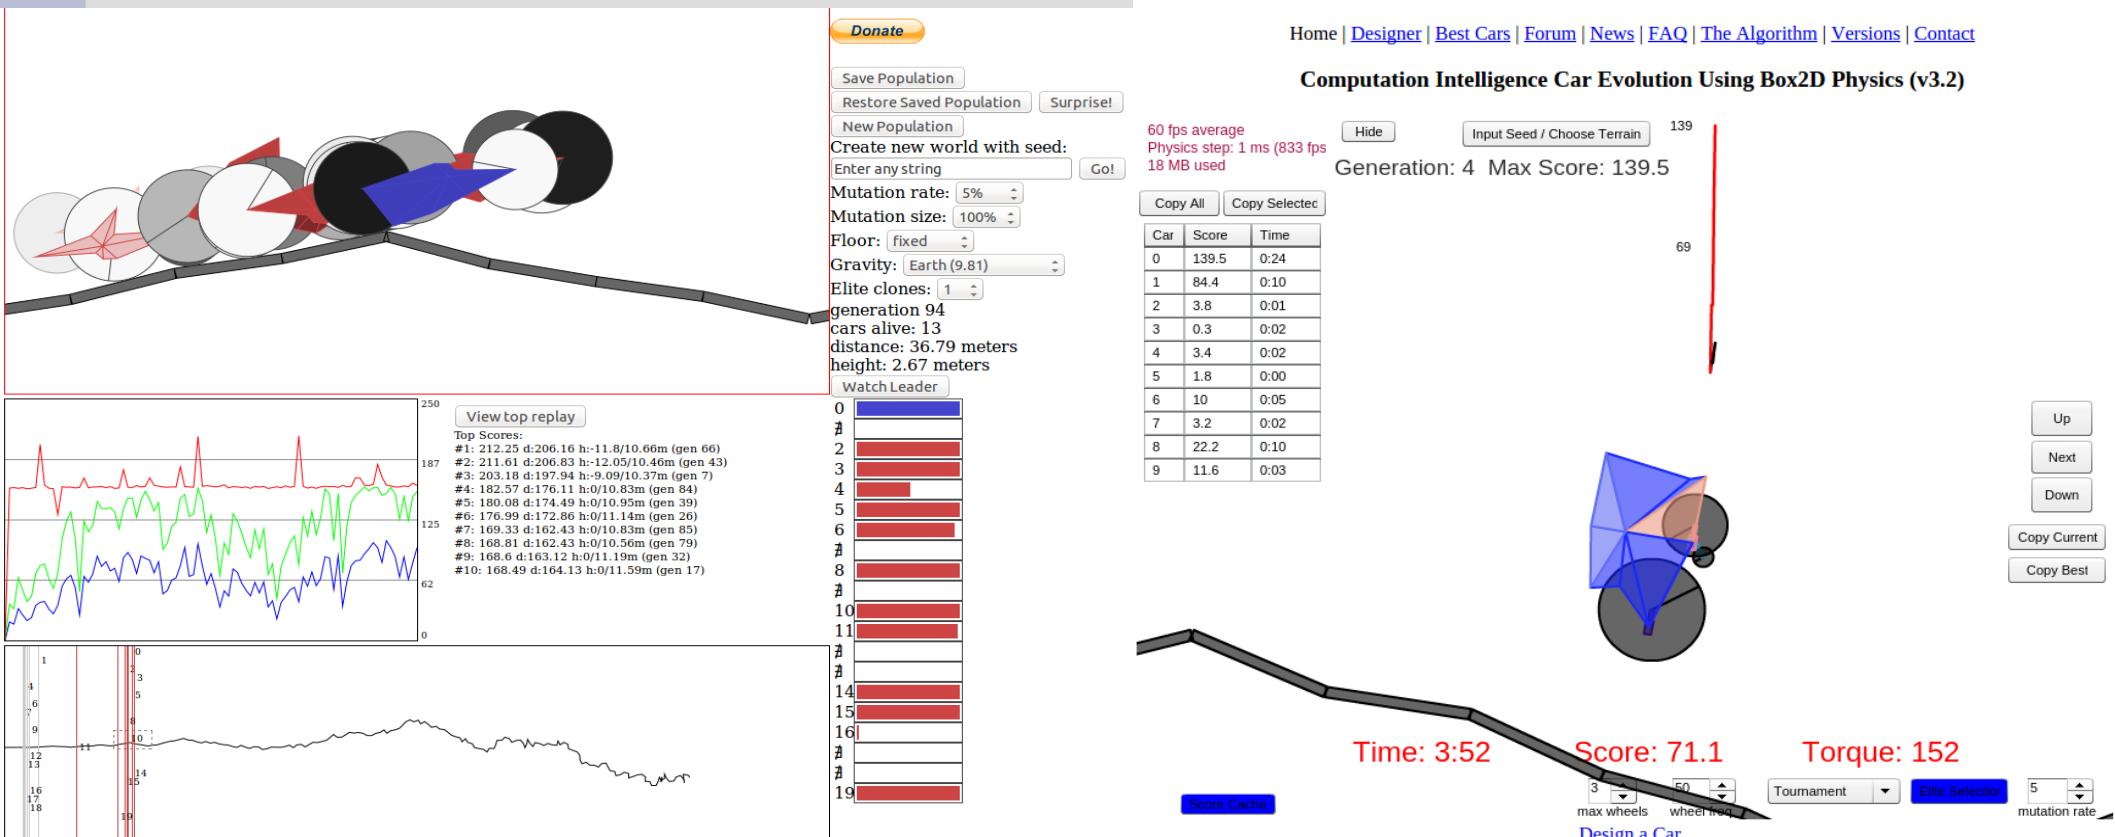
Top Scores:
#1: 212.25 d:206.16 h:-11.8/10.66m (gen 66)
#2: 211.61 d:206.83 h:-12.05/10.46m (gen 43)
#3: 203.18 d:197.94 h:-9.09/10.37m (gen 7)
#4: 182.57 d:176.11 h:0/10.83m (gen 84)
#5: 180.08 d:174.49 h:0/10.95m (gen 39)
#6: 176.99 d:172.86 h:0/11.14m (gen 26)
#7: 169.33 d:162.43 h:0/10.83m (gen 85)
#8: 168.81 d:162.43 h:0/10.56m (gen 79)
#9: 168.6 d:163.12 h:0/11.19m (gen 32)
#10: 168.49 d:164.13 h:0/11.59m (gen 17)

Score Cache

Time: 3:52 Score: 71.1 Torque: 152

max wheels: 3 wheel freq: 50 Tournament Elite Selector mutation rate: 5

Design a Car



Genetic algorithms

Genetic algorithms are based on how life has evolved over time

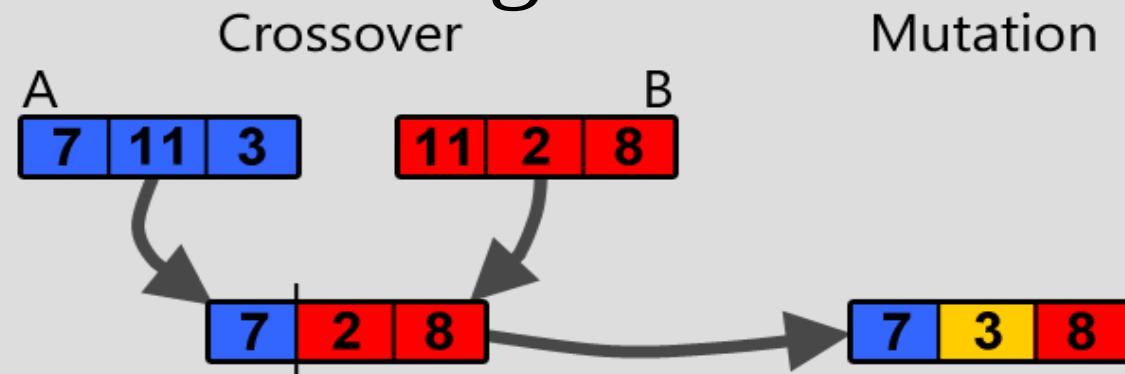
They (in general) have 3 (or 5) parts:

1. Select/generate children
 - 1a. Select 2 random parents
 - 1b. Mutate/crossover
2. Test fitness of children to see if they survive
3. Repeat until convergence

Genetic algorithms

Selection/survival:

Typically children have a probabilistic survival rate (randomness ensures genetic diversity)



Crossover:

Split the parent's information into two parts, then take part 1 from parent A and 2 from B

Mutation:

Change a random part to a random value

Genetic algorithms

Genetic algorithms are very good at optimizing the fitness evaluation function (assuming fitness fairly continuous)

While you have to choose parameters (i.e. mutation frequency, how often to take a gene, etc.), typically GAs converge for most

The downside is that often it takes many generations to converge to the optimal

Genetic algorithms

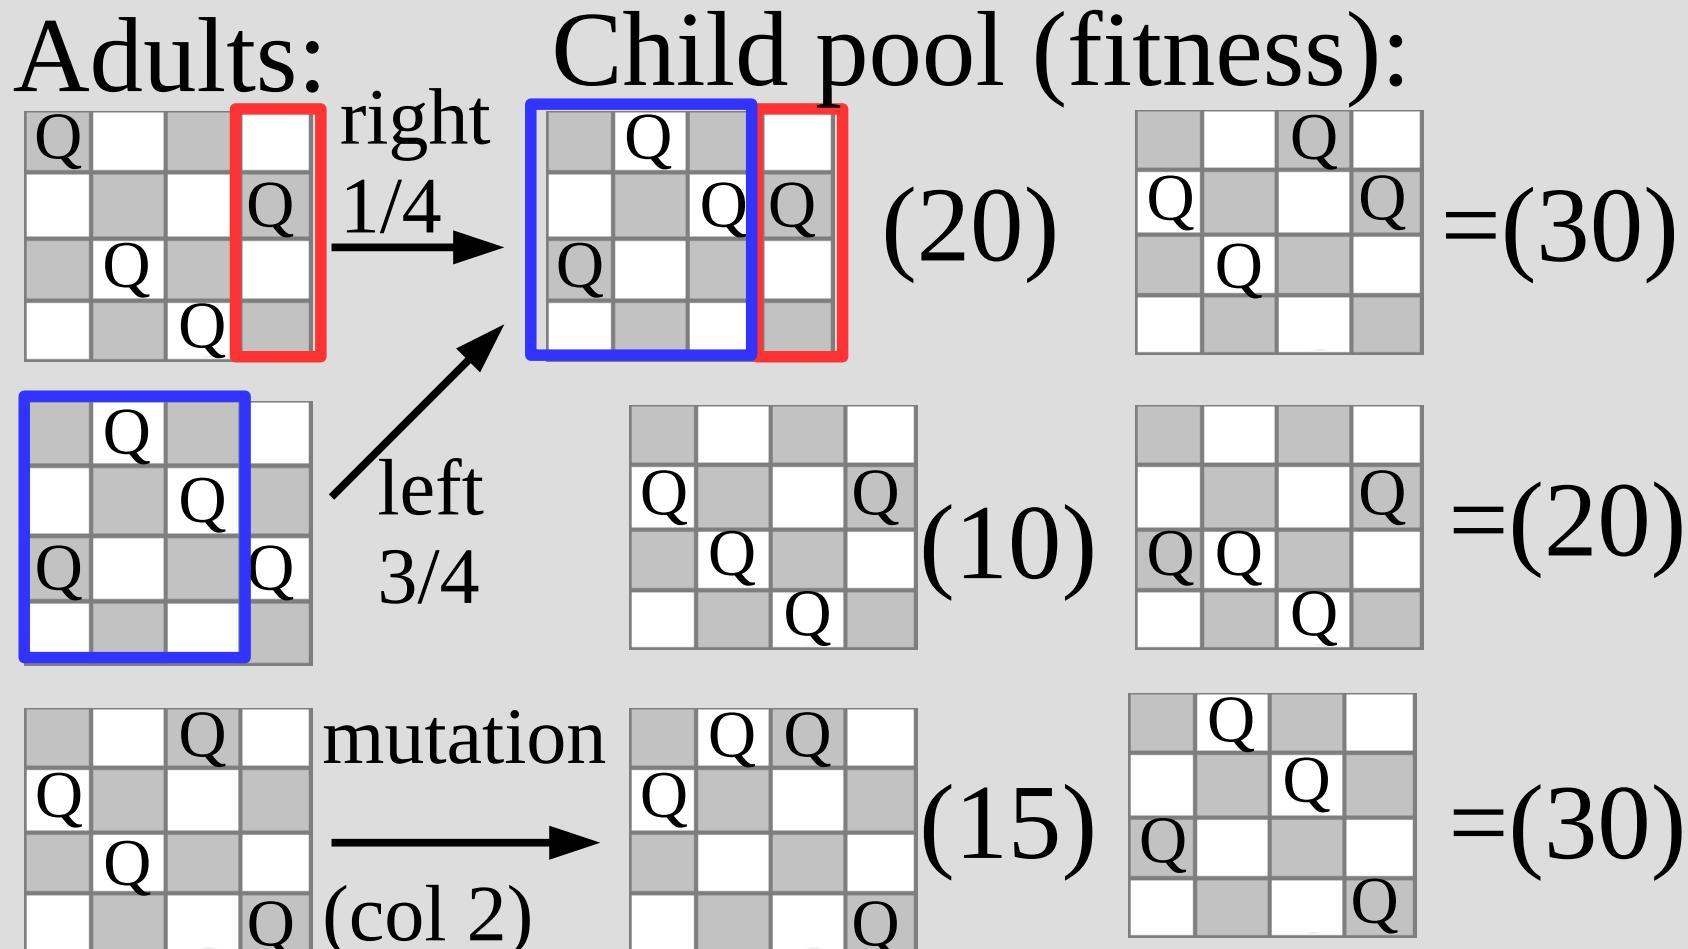
There are a wide range of options for selecting who to bring to the next generation:

- always the top (similar to hill-climbing... gets stuck a lot)
- choose purely by weighted random (i.e. 4 fitness chosen twice as much as 2 fitness)
- choose the best and others weighted random

Can get stuck if pool's diversity becomes too little (hope for many random mutations)

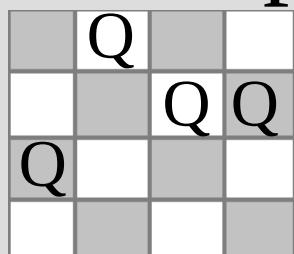
Genetic algorithms

Let's make a small (fake) example with the 4-queens problem

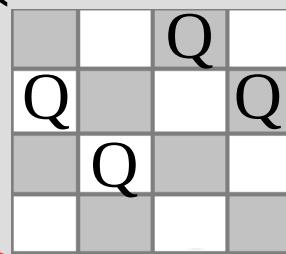


Genetic algorithms

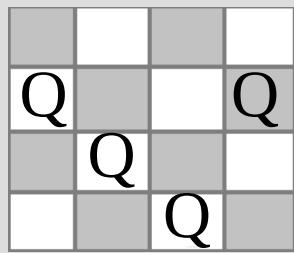
Let's make a small (fake) example with the
4-queens problem
Child pool (fitness):



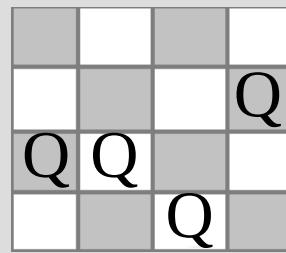
(20)



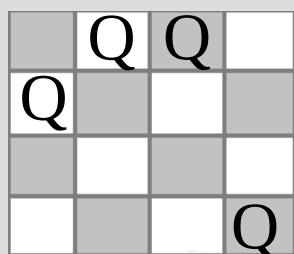
= (30)



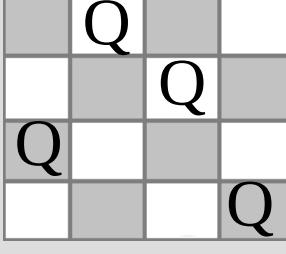
(10)



= (20)

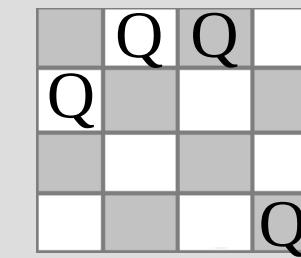
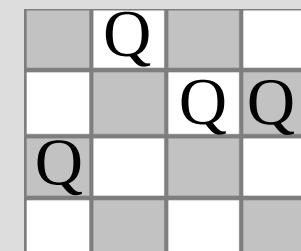
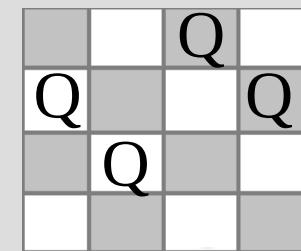


(15)

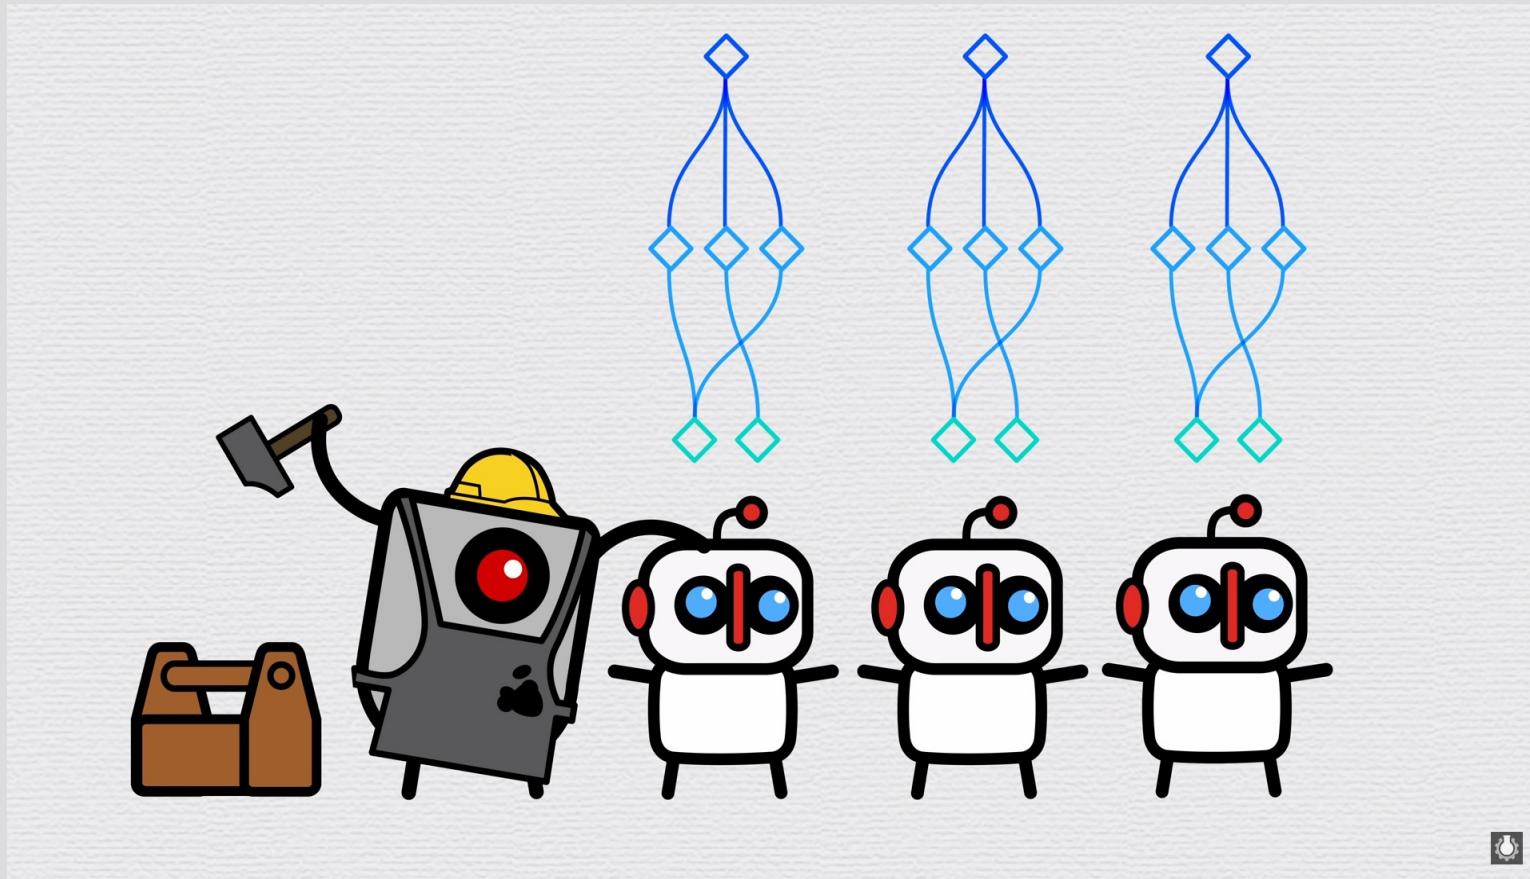


= (35)

Weighted random
selection:



Genetic algorithms



<https://www.youtube.com/watch?v=R9OHn5ZF4Uo>