

Modern Robotics: Evolutionary Robotics COSC 4560 / COSC 5560

Professor Cheney 1/29/18

Variations of Evolutionary Algorithms

Fitness/Search Landscapes











but wouldn't we rather have a global optima?

to do so, we'd have to accept (many) negative mutations to get to the better "fitness peak"



what if just sometimes, we accepted negative mutations?

stochastic hillclimbing



what if instead we used random initialization?









how do you decide the probability of (when to) accept a negative mutation or not?

(i.e. when should you explore and when should you exploit?)

explore (take risky moves) early in search so that you have more time to catch up (and exploit) later on

simulated annealing

the probability of accepting a new negative mutation decreases over optimization time

at first the point (current state) is randomly moving around all over the place (state space)

but as the system "temperature cools" over time, it settles down, and resists change and will only accept a new position if it's an improvement over the current one





genetic algorithms

population based methods







Selection Pressure

How much "pressure" is there on a solution to make an uphill move ("exploitation") vs. a downhill move ("exploration")? As we saw from random-restart hillclimbing and simulated annealing, we want to explore earlier in optimization (to find the basin of our global optima), and exploit later in optimization (to climb to the optima of that basin)

Easier said than done...

Truncation Selection

Select the top-N individuals



Maximum selection pressure (always moving uphill!)

Fitness-Proportional Selection

Select N individuals, with an individual being selected at each draw by the proportion of fitness it represents (often without replacement)

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \\ 30 \end{aligned}$ 7/30 = 23%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 30 \end{aligned}$ 7/30 = 23%

Select N individuals, with an individual being selected at each draw by the proportion of fitness it represents (often without replacement)

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0$

This is great for stochastically favoring better individuals more the higher their fitness is

But it comes with one major flaw too...

Suppose that later in optimization, fitness values have gone up dramatically

Selection pressures is now much lower (i.e. smaller chance of taking an uphill step) than early on when average fitness values were smaller (this is the opposite of what we want!)

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}$)= 107107/630=17.0%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}$)= 106106/630=16.8%Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$)= 105105/630=16.7%Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%

Selection pressures is now much lower (i.e. smaller chance of taking an uphill step) than early on when average fitness values were smaller (this is the opposite of what we want!)

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}$)= 107107/630=17.0%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}$)= 106106/630=16.8%Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$)= 105105/630=16.7%Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$)= 104104/630=16.5%

Rank-based Selection

Define some selection proportion for each rank in the population *a priori*

e.g. exponential fall-off

1) 50%
 2) 25%
 3) 12.5%
 4) 6.25%
 5) 3.13%
 6) 1.56%

We just apply are preset selection probabilities

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix} = 107$ 50%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} = 106$ 25%Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} = 105$ 12.5%Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix} = 104$ 6.25%Fit($\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix} = 104$ 3.13%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix} = 104$ 1.56%

And selection pressure is constant as average fitness changes

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix} = 7$ 50%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} = 6$ 25%Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} = 5$ 12.5%Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix} = 4$ 6.25%Fit($\begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \end{bmatrix} = 4$ 3.13%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix} = 4$ 1.56%

Often we always want to include the most fit individual in the next population ("elitism")

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix} = 7$ 100%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} = 6$ 25%Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} = 5$ 12.5%Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix} = 4$ 6.25%Fit($\begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \end{bmatrix} = 4$ 3.13%Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix} = 4$ 1.56%

Tournament Selection

Randomly select a subset (mini-batch) of your population, perform some other selection method (most frequently truncation selection) only on that subset and return the winners Tournaments can be of any size

e.g. tournament of size 2

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}$) = 7 Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix}$) = 6 Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$) = 5 Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$) = 4 Fit($\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}$) = 4 Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$) = 4 Tournaments can be of any size

e.g. tournament of size 2

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}$) = 7 Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}$) = 6 Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$) = 5 Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$) = 4 Fit($\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$) = 4 Usually you continue to select additional tournaments

→ Fit([1 0 1 0 1 1 1 0 1 1]) = 7 ✓
Fit([1 1 1 0 1 0 1 0 1 0 1]) = 6 ✓
Fit([0 1 1 0 1 0 1 0 0 1 0]) = 5
Fit([0 0 1 1 0 1 0 1 0 1 0]) = 4
Fit([0 0 0 0 1 0 1 0 1 1 0]) = 4

Usually you continue to select additional tournaments until all individuals have been evaluated and selected

Small tournaments (e.g. size 2), have low section pressure since highly fit individuals can be knocked out

Small tournaments (e.g. size 2), have low section pressure since low fit individuals can survive

Fit($\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}$) = 7 Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}$) = 6 Fit($\begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$) = 5 Fit($\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$) = 4 Fit($\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$) = 4 Fit($\begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$) = 4 At the other extreme, a tournament of the full populations is just regular truncation selection (and thus has maximum selection pressure)



So tournament size is a great knob to use for controlling selection pressure

Tournaments are also great in that they can be used in "steady-state" algorithms, that perform variation-evaluation-selection loops on small subsets of the population at a time, instead of "generational" algorithms that have distinct iterations