

Evolutionary Robotics:



Genetic algorithms (Holland 1975) have been put to a wide variety of search and optimisation tasks from genetic programming (Koza 1992) to evolutionary robotics (Husbands et al 1996).

Artificial evolution begins with a population of strings or genotypes describing the control parameters of an agent. Each agent is evaluated in the problem domain and given a measure of fitness.

Following the inspiration from Darwinian evolution (Darwin 1759) the fitter genotypes have a higher probability of breeding.

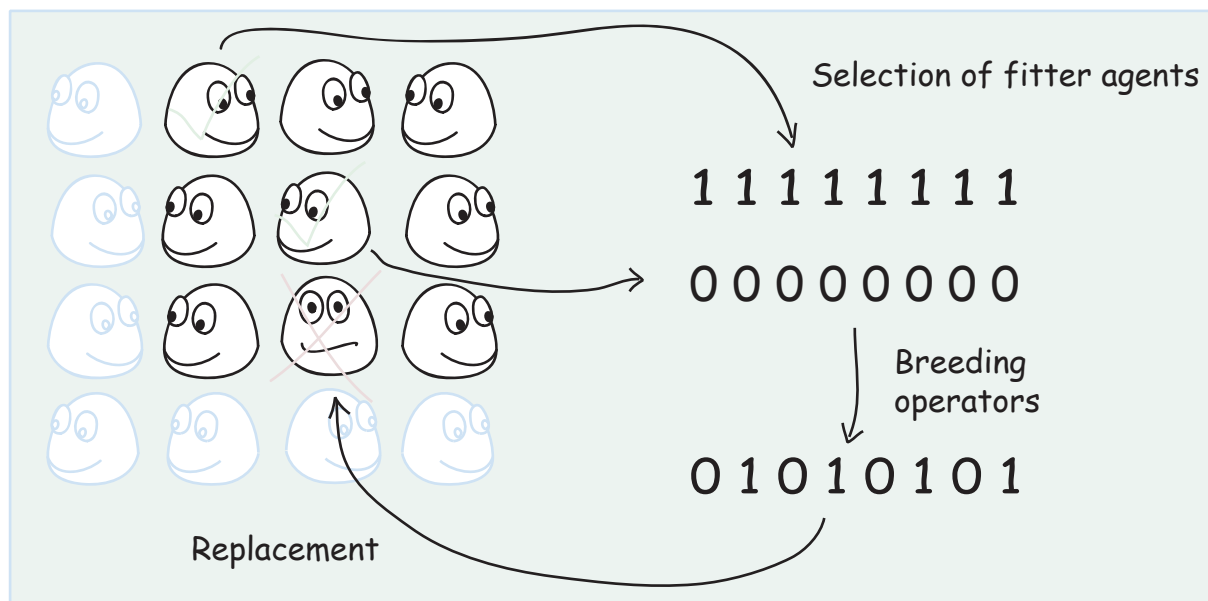
Breeding usually involves crossing two parents to produce a child and modifying or 'mutating' occasional values.

As the generations pass, mean fitness within the task domain increases.



The spatially distributed genetic algorithm:

For each iteration an agent is selected at random from a population distributed across a square grid. The agent is bred with one of its fittest neighbours and their offspring replaces one of the least fit neighbours.



□ The selection pressure keeps successful genes in the population while breeding operators facilitate exploration of the neutral space around the population.

□ The distributed GA is a fast and robust optimisation procedure. After just a few generations capable agents appear.