Evolutionary Robotics
Prof. Dario Floreano

Evolutionary Robotics Laboratory

Exercise Sheet 0: Evolutionary Algorithms – Responses

Euan Judd (euan.judd@epfl.ch)
Krishna Manaswi Digumarti (krishna.digumarti@epfl.ch)

Exercise 1

Most mutations are deleterious: they worsen fitness, but some will improve the fitness of the parent.

If the parent is very far from the optimum then large mutations could rapidly move the individuals closer to the optimum. However, if the parent is already close to the optimum then large mutations will bring us further away. Small mutations are therefore needed to "hone in" on the optimum. Since the sphere function is separable (i.e. it is possible to optimize each dimension in isolation) the dimensionality should not influence the mutation rate.

Exercise 2

An intermediate mutation rate tends to do the best because it offers a compromise between the ability to make large jumps when the population is far away from the optimum and small changes to hone in once the population approaches the optimum. Later in the class you will see more sophisticated methods that dynamically adapt the mutation rates instead of always using a fixed value.

Exercise 3

Crossover alone does not work well because the only genetic material that will ever exist is that which was present in the initial population. Mutation is needed to introduce novel genetic material.

Since the sphere function in separable (see above) always using crossover (plus some mutation) is a good idea. As a small example, consider the 2D case where we mate individuals [0, 0.5] and [0.9, 0]. Crossover could cause us to jump directly to an offspring with the optimum [0, 0] if the first and second indices were selected from the first and second individuals, respectively.

Exercise 4

The first thing to realize is that the tournament size influences the selection pressure (the larger the tournament, the greater the selection pressure). Greater selection pressure means we are doing more "exploitation" than "exploration". Since the sphere function is unimodal, any fitness improvement takes us closer to the global optimum, so it is good to have strong selection pressure and therefore a large tournament size. On the other hand, the Rastrigin function is multimodal (there are many local optima). In this case if we have too strong selection pressure (too large a tournament) it is likely that the population will become stuck in one of the local optima. So, using a small tournament size is better because it allows more exploration and therefore a greater chance of finding the global optimum.

Exercise 5

This exercise should have given you a further opportunity to see how the problem size (i.e. its dimensionality/number of variables), as well as the presence of many local optima, influences how difficult it is for an EA to solve.