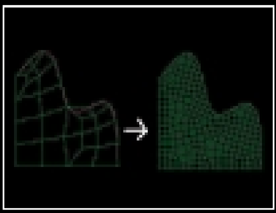

GA Optimization of Approximate or Noisy Function Representations

Evolutionary Optimization Group

The optimization of noisy or imprecisely specified functions is a common problem occurring in various applications. In many optimization problems there may exist a number of different ways in which a particular problem is modelled. Some methods may be quite elaborate in their representation, while others involve a simplification of the problem, with the former being more accurate but at the same time more computationally expensive than the latter. Typical examples from engineering might be a coarse FEA mesh as compared to a refined one for stress analysis, or CFD panel and Euler methods for estimating drag.

It is therefore important to understand how a significant number of less accurate evaluations could be integrated with fewer accurate ones to arrive at an optimum design. To understand how different integration strategies work, we used the generalized bump function as detailed below.

The α and β parameters are used here to define frequency and spatial distortion, respectively. The undistorted function is one in which α=1 and β=0.

An interesting feature of this function is that the surface is nearly but not quite symmetrical in $x_1=x_2$, so the peaks always occur in pairs but one is always bigger than its sibling.

The global optimum is defined by the product constraint. When the problem is generalized for n greater than two it becomes even more demanding with families of similar peaks occurring within a highly complex constraint surface. Shown below is the 2D undistorted bump (golden) with a distorted variant superimposed (white translucent). Notice the many false and distorted optima.

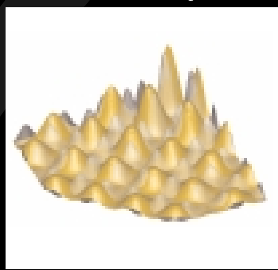

Two GA optimisation methods have been used on this problem. One involving a simple haploid GA, and the other a GA with niching. Below is the distribution of the population over the function for the simple GA. It is clear that most of the population has converged on a single peak which is not the global maximum (the global maximum lies on the constraint boundary).

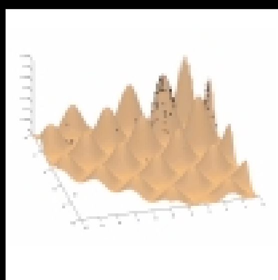

On the other hand a GA with niching forces the population to be distributed over many peaks in the parameter space as is seen below

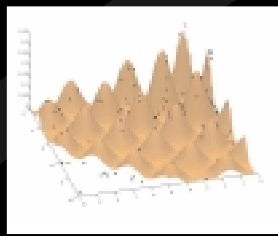

Both types of GAs have been applied to this problem with a changing fitness landscape. The optimisation was carried out sequentially over three levels starting with the highly distorted function, then a less distorted one, and ending with the base function. More generations were allocated on the highly distorted function (analogous to a coarse representation that is computationally cheap), fewer were allocated to the less distorted bump, and even fewer to the base function. This process has been carried out for both frequency distortion (varying α) and spatial shift (varying β).

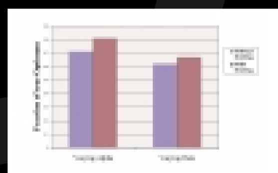

Results for this study have shown that the use of niching, which distributes the population over many peaks in a changing fitness landscape, gives improved results.

Further work is to be carried out on higher dimension functions and also by implementing diploid GAs that may be better suited for such changing landscapes.

Mohammed El-Beltagy
Evolutionary Optimization Group
Computational Engineering and Design Centre
University of Southampton
Highfield, Southampton
SO17 1BJ
Email: M.A.El-Beltagy@soton.ac.uk
Tel +44-1703-593392
Fax +44-1703-593230.

University of Southampton