

experiments carried out which compare the performance on a range of $C(N, K)$ landscapes of three crossover operators (namely one-point, two-point and uniform) followed by a discussion of the results and how each operator differs in the nature of its exploration of the search space.

3 An NK Model of Fitness Landscapes for Real-Valued Strings

In the context of NK landscapes, the use of a lookup table is impracticable when generating fitness functions for real-valued, as opposed to binary, strings since the number of allelic combinations is infinite. The manner in which the fitness contribution of each locus is calculated has been modified as follows: instead of being compared with the entries of a lookup table, the alleles of the current locus and its interactants are summed and the result multiplied by an arbitrary weight, ω_{x_i} , associated with that locus; this number is then used as the seed, t_{x_i} , for a pseudo-random number generator, Ψ . The first pseudo-random number generated using this seed is taken as the fitness contribution of the locus. (The numbers generated are uniformly distributed over the interval $[0.0, 1.0]$.) More formally,

$$F_i(x_i) = \Psi_1(t_{x_i}) \quad (2)$$

where

$$t_{x_i} = \omega_{x_i} \left(x_i + \sum_{k=1}^K x_{i_k} \right) \quad (3)$$

(In this case, $0 \leq x_i \leq 100$ and $1 \leq \omega \leq 10,000$; however, these choices were arbitrary.) This method of combining the interactants' alleles to form what is likely to be a unique seed, and assigning the first number generated using that seed to a fitness contribution, retains the random nature of the assignment while ensuring that a particular allelic combination will always make the same contribution. $C(N, K)$ landscapes are therefore analagous to Kauffman's original NK landscapes.

3.1 Continuous NKp Fitness Landscapes

The NK model assumes that every combination of interactants makes a positive contribution to fitness; in nature (and real-world problems?), it seems plausible that many combinations of alleles will make no contribution to fitness [1]. This can be reflected in the NK model by assigning such a combination 0 in the fitness table. Hence, an NKp landscape is one in which each allelic combination is assigned a fitness contribution equal to 0 with probability p ; the case $p = 0$ corresponds to a normal NK landscape, while $p = 1$ results in all fitness table entries being set to 0, and therefore a completely flat landscape.

The method of constructing $C(N, K)$ landscapes described above in section 3 was modified to implement $C(N, K, p)$ landscapes quite simply as follows: recall that the fitness contribution, F_i , of a locus is a pseudo-randomly generated number between 0 and 1; this number is now passed to the probability function used to decide whether or not the allelic combination in question is to make a fitness contribution. Moreover, if $p \geq F_i$ then F_i is set to 0.

3.2 Experiments

Experiments were carried out for $N = 100$ and $K = 0, 5, 20, 50$ and 80 . An oft-cited shortcoming of NK landscapes is that all genes are subject to the same degree of epistasis (since traditionally K is fixed), while real-world problems vary considerably in the amount of epistasis between parameter subsets [4]. Additional experiments were therefore carried out for variable K : when constructing a landscape each locus was assigned a random number of interactants between 0 and an upper value for K ; namely, $K \leq 30, 60, 0$. Random and adjacent linkage types were used for all values of K ; 50 runs of 500 generations were carried out for all values of K , each randomly generating a new landscape at its outset.

The crossover operators compared were one-point (1PX), two-point (2PX), and uniform (UX) crossovers.



3.3 The Genetic Algorithm

A geographically distributed GA was used, with the population spread across a two-dimensional toroidal grid of size 15×15 , each cell of which contained a single individual. Mating was restricted to small groups of individuals which were generated as follows:

<i>K</i>	1PX			2PX			UX		
	Best	μ	σ	Best	μ	σ	Best	μ	σ
0	0.854	0.810	0.02	0.888	0.855	0.013	0.57	0.4	0.004
5R	0.724	0.671	0.01	0.741	0.62	0.015	0.820	0.788	0.013
5A	0.772	0.748	0.013	0.75	0.771	0.012	0.808	0.781	0.013
20R	0.656	0.626	0.013	0.671	0.636	0.013	0.707	0.676	0.012
20A	0.710	0.678	0.012						

K	1PX			2PX			UX		
	Best	μ	σ	Best	μ	σ	Best	μ	σ
$\leq 30R$	0.686	0.651	0.014	0.706	0.665	0.016	0.806	0.753	0.018
$\leq 30A$	0.750	0.706	0.016	0.753	0.718	0.014	0.772	0.750	0.014
$\leq 60R$	0.665	0.633	0.012	0.67	0.641	0.014	0.751	0.708	0.015
$\leq 60A$	0.704	0.675	0.015	0.717	0.686	0.013	0.742	0.711	0.014
$\leq 0R$	0.667	0.625	0.012	0.66	0.634	0.011	0.735	0.688	0.017
$\leq 0A$	0.68	0.660	0.013	0.704	0.668	0.014	0.737	0.6 0	0.016

Table 2: Results for variable K






Figure 2: Comparison of operators’ distributional bias averaged over 50 runs.

K	1PX	2PX	UX
0	1	26	52
5R		12	52
5A	12	16	52
20R	8	18	52
20A	8	11	54
50R	10	23	52
50A	6	12	51
80R	10	23	51
80A	10	21	53

Table 3: Number of contributing individuals to final solutions, averaged over 50 runs

Interpretation of positional bias can be misleading; as stated above, UX and 2PX exhibit the same positional bias (that is, that all genes are equally likely to be exchanged); however, this disregards the probabilities of certain genes being exchanged *together* – probabilities which, quite obviously, will differ significantly between the two operators.

5.3 Step Size

Typical examples of average step-size and their standard deviation are illustrated in figure 4. As with the two biases, the graphs of step sizes taken over the course of a run reflect the crossover operators’ rates of convergence; however, it can be seen that, during the early generations, while UX takes significantly larger steps on average than either 1PX or 2PX takes, the size of its steps varies much less than those two operators. This suggests that UX effects were wider and more uniform coverage of the fitness landscape (and therefore a more efficient global search) than do 1PX and 2PX. For 1PX and 2PX, the decrease of step size (like the biases) slows as K increases, and adjacent linkage causes decreases which are faster than random linkage for equivalent values of K .

6 Discussion

The one- and two-point operators are clearly sensitive to the nature, in addition to the amount, of epistasis present in a problem, as illustrated by the contrast in behaviour between adjacent and random linkage. Patterns observed in positional bias, distributional bias and step-size are all characterized by that of convergence; therefore, it is perhaps the rate of convergence which provides an explanation for the operators’ aforementioned epistatic sensitivity.

The rate of convergence for both 1PX and 2PX is slower when performance is lower, namely when linkage between interacting genes is random. Convergence – premature or not – occurs because (segments of) genes are repeatedly propagated throughout the population; the more the same genes are propagated

