Forma Analysis and Random Respectful Recombination

Nicholas J. Radcliffe

njr@castle.ed.ac.uk

Edinburgh Parallel Computing Centre, King's Buildings, University of Edinburgh, EH9 3JZ, Scotland

Abstract

Intrinsic parallelism is shown to have application beyond schemata and *o*-schemata. More general objects called *formae* are introduced and general operators which manipulate these are introduced and discussed. These include *random, respectful recombination.* The extended formalism is applied to various common representations and standard operators are analysed in the light of the formalism.

1 Introduction

The conventional understanding of genetic algorithms attributes much of their power to intrinsic parallelism, the phenomenon whereby each chromosome is an instance of many schemata (or *o*-schemata) and its measured performance contributes to an estimated fitness for each of these schemata. Efforts to maximise the level of intrinsic parallelism available are frequently in conflict with a desire to use natural representations and operators for the structures in space being searched. This paper demonstrates that intrinsic parallelism is a very general phenomenon, not restricted to schemata and *o*-schemata and explores the interaction between intrinsic parallelism, genetic representations and operators.

The paper begins (sections 2 & 3) with a review of earlier work which showed that more general partitions of the search space than schemata give rise to intrinsic parallelism. This motivates a shift of emphasis from schemata to more general kinds of regularities which may be present in the search space (sections 4 & 6), and allows the introduction of general-purpose operators—including so-called *random, respectful recombination*—which can be of both analytic and practical use (section 5). The generalisation of a schema is called a *forma*, and in sections 7 to 10 four different types of formae[1] are discussed, in conjunction with operators for their effective manipulation. Before concluding, there is a discussion of the future directions suggested by this work.

[1] Although Holland chose the neuter form for the Latin noun schema, there is no option but to choose the feminine form of its synonym, forma.

2 Analysis of Genetic Algorithms

In his seminal book on adaptation, John Holland [7] considered chromosomes which were k-ary[2] strings or similar and showed that genetic algorithms can usefully be analysed in terms of their effects on higher-order structures, variously known as *schemata, hyperplanes* or *similarity templates.* A schema specifies some set of alleles which the genes of a chromosome must express in order for that chromosome to be said to be an *instance* of the schema. Holland derived a simple but immensely significant expression which bounded the expected number of instances of any schema in the next generation of the population, commonly known as the "schema" or "fundamental" theorem. Later, David Goldberg [5, 4] introduced a number of classes of "*o*-schemata" which played much the same rôle for problems in which the chromosome was a member of the permutation group $\mathcal{P}_n$ (such as the travelling sales-rep problem, TSP) as did Holland's original schemata for k-ary string representations. The following analysis shows that schemata and *o*-schemata can both be regarded as examples of a more general kind of object which we shall term a *forma.*

For present purposes it will be convenient to identify a forma (currently a schema or an *o*-schema) with the set of all of its instances, so that if a chromosome η is an instance of a forma ξ we shall write $\eta \in \xi$. In this spirit, $3241 \in 3\square 4\square$ where $\square$ is the "don't care" symbol.

Recall that the Fundamental Theorem bounds the expected number of instances $N_\xi(t+1)$ of each forma ξ in the population $\mathfrak{B}(t+1)$ by

$$\left\langle N_\xi(t+1) \right\rangle \geq N_\xi(t) \frac{\hat{\mu}_\xi(t)}{\bar{\mu}(t)} \left[1 - \sum_{\omega \in \Omega} p_\omega p_\omega^\xi \right],$$

where $\hat{\mu}_\xi(t)$ is the sample average for utility of ξ over all its instances in the population $\mathfrak{B}(t)$ and the terms $p_\omega p_\omega^\xi$ in the sum quantify the disruptive effect of each operator ω, drawn from a set Ω of genetic operators, on forma membership.

Recall also that the *defining positions* of a forma are those loci at which a value is specified, so that the forma $\xi = \square a \square b$, has defining positions at its second and fourth loci.

[2] base k, for arbitrary k

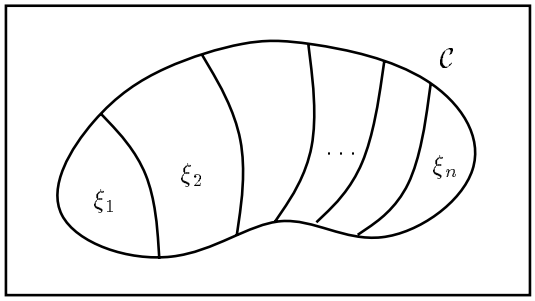

Figure 1: An equivalence relation partitions the space of chromosomes, $\mathcal{C}$, into a number of equivalence classes or *formae*, $\xi_1, \xi_2, \ldots, \xi_n$.

The generalisation sought in this paper requires us to introduce equivalence relations over the search space. Given any forma ξ, consider the equivalence relation which relates any pair of chromosomes having the same alleles at the forma's defining positions. We can choose to regard the forma as an equivalence class induced by this equivalence relation. Specifically, we can denote the equivalence relation which induces $\square a \square b$ by $\square \blacksquare \square \blacksquare$, which we understand to relate those chromosomes which have the same values at those loci marked with the $\blacksquare$ symbol, placing 1234 and 2224 in the forma $\square 2 \square 4$ but 1111 and 0101 in the forma $\square 1 \square 1$. We shall see examples of formae which are less like familiar schemata in section 10.

Having made this identification, the Fundamental Theorem can be seen to apply to *any* subset ξ of the space of chromosomes, $\mathcal{C}$, provided only that the disruption coefficient p_ω^ξ correctly bounds the disruptive effect of applying the operator ω to a chromosome which is a member of ξ. In practice we shall choose to regard these subsets as equivalence classes induced by some set of equivalence relations—a freedom we always have. For this reason we shall henceforth use the term "forma" to refer to an equivalence class of any equivalence relation over the space $\mathcal{C}$ of chromosomes. Holland's schemata and Goldberg's various *o*-schemata are then immediately seen to be special cases of formae. (See figure 1.)

3 Exploiting Intrinsic Parallelism

Holland observed that each evaluation of a chromosome can be regarded as a statistical sampling event which yields information about the sample averages for utility of *each* of the 2^n schemata of which it is an instance (the phenomenon referred to as intrinsic parallelism) but of course this applies equally to any formae we may choose to consider. The parallelism is exploited by using information gathered about these higher order structures, the formae, to guide the further exploration of the space. The critical tasks are thus finding formae which characterise solutions in meaningful ways and developing operators which usefully manipulate these formae.

It is important to notice that the primary factor governing the expected rate of increase of (instances of) any forma ξ is not the mean relative fitness of its members, $\mu_\xi / \bar{\mu}(t)$, but the *observed* relative fitness $\hat{\mu}_\xi(t)/\bar{\mu}(t)$ of its instances in the population $\mathfrak{B}(t)$. For this reason, our effective exploitation of information about the fitness of various formae is strictly limited by the reliability of the sample fitness $\hat{\mu}_\xi(t)$ as an estimator of the mean fitness μ_ξ of *all* instances of ξ. This suggests the unsurprising conclusion that we will only be able to exploit effectively information about formae whose instances display a low variance for fitness. More succinctly, only those formae which well-characterise solutions, identifying sets with broadly similar performance, will be of any value to the search. Thus the degree of intrinsic parallelism which can be effectively utilised by the search is limited to the number of formae which capture regularities in the performance of solutions over the space $\mathcal{C}$ of chromosomes.

These considerations and others (including Goldberg's principles of *minimal alphabets* and *meaningful building blocks*, [4]) led to the proposal of six *design principles* for constructing useful representations, formae and genetic operators (Radcliffe [10]). In the following, the number of formae induced by an equivalence relation will be referred to as the *precision* of both the relation and the formae it induces.[3] The set of equivalence relations which induce the formae (equivalence classes) under consideration will be written Ψ and the set of all formae induced by relations in Ψ will be denoted Ξ.

Two formae $\xi, \xi' \in \Xi$ will be said to be *compatible* if it is possible for a chromosome to be an instance of both ξ and ξ' (figure 2). In the familiar case of schemata, $1\square 1\square$ and $0\square\square\square$ are incompatible, because there is a conflict at the first locus, whereas $1\square 1\square$ and $11\square\square$ are compatible.

In general, recombination operators take two chromosomes and produce different children depending on explicit or implicit control parameters such as the crosspoint used for one-point crossover and the binary mask used in uniform crossover. (See Syswerda [11] and Eshelman *et al* [3] for details of uniform crossover.) A generic recombination operator X will be taken to have an associated control set $\mathcal{A}_X$ and functional form

$$X : \mathcal{C} \times \mathcal{C} \times \mathcal{A}_X \longrightarrow \mathcal{C}.$$

The member of this control set chosen for some particular recombination completely determines which of the various

[3] In the case of schemata and genes with k alleles, the precision is $k^{\,o}$, where o is the order of a schema.

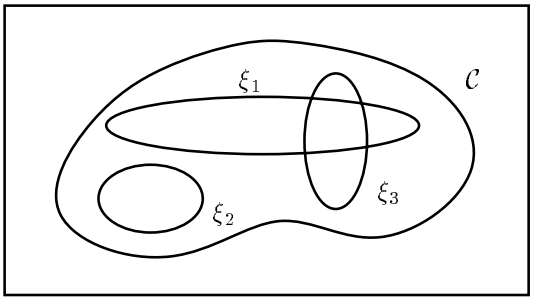

Figure 2: The formae ξ_1 and ξ_2 are *incompatible* because they have null intersection, whereas ξ_1 and ξ_3 are compatible because a single chromosome can be an instance of them both.

possible children results from the cross.

4 Design Principles

The first three design principles are general and suggest desirable characteristics of chromosomal representations and formae:

1. (Minimal redundancy) *The representation should have minimal redundancy; such redundancy as exists should be capable of being expressed in terms of the equivalence relations in* Ψ.
Ideally, each member of the space being searched should be represented by only one chromosome in $\mathcal{C}$. This is highly desirable in order to minimise the size of the search space. If redundant solutions are related by one of the equivalence relations used then the genetic algorithm should effectively be able to "fold out" the redundancy (see principle 4); otherwise it is doomed to treat redundant solutions as unrelated.

2. (Correlation within formae) *Some of the equivalence relations, including some of low precision, must relate chromosomes with correlated performance.*
This ensures that useful information can be gathered about the performance of a forma by sampling its instances. Such information is used to guide the search. The emphasis is placed on low-precision formae because these will generally be less likely to be disrupted by the application of genetic operators, and are also more likely to be compatible with one another.

3. (Closure) *The intersection of any pair of compatible formae should itself be a forma.*
This ensures that solutions can be specified with different degrees of accuracy and allows the search gradually to be refined. Clearly the precision of formae so constructed will be at least as high as that of the higher-precision of the intersecting formae.

The remaining three principles concern the way in which operators manipulate chromosomes and formae. It is helpful to use an informal analogy in which chromosomes specify people and some of the characteristics used to define a set of formae are hair colour and eye colour. The consequence of each design principle for these formae is given at the end of each principle.

4. (Respect) *Crossing two instances of any forma should produce another instance of that forma.*
Formally, it should be the case that

$$\forall \xi \in \Xi \ \forall \eta, \zeta \in \xi \ \forall a \in \mathcal{A}_X \ : \ X(\eta, \zeta, a) \in \xi,$$

where X is the crossover operator. In this case the crossover operator will be said to *respect* the equivalence relations (and their formae). This is necessary in order that the algorithm can converge on good formae, and implies, for example, that $X(\eta, \eta, a) \equiv \eta$, assuming that equivalence relations of maximum precision specify chromosomes completely.
[If both parents have blue eyes then all their children produced by recombination must have blue eyes.]

5. (Proper assortment) *Given instances of two compatible formae, it should be possible to cross them to produce a child which is an instance of both formae.*
Formally,

$$\forall \xi, \xi' \in \Xi \ (\xi \cap \xi' \neq \emptyset \ ) \ \forall \eta \in \xi \ \forall \eta' \in \xi' \\ \exists \, a \in \mathcal{A}_X \ : \ X(\eta, \eta', a) \in \xi \cap \xi'. \qquad (1)$$

This relates to Goldberg's "meaningful building blocks", of which he writes ([4], p. 373)

> 'Effective processing by genetic algorithms occurs when *building blocks*—relatively short, low order schemata with above average fitness values—combine to form optima or near-optima.'

A crossover operator which obeys equation 1 seems very much more likely to be able to recombine "building blocks" usefully, and any crossover operator which obeys this principle will be said *properly to assort* formae.
[If one parent has blue eyes and the other has brown hair it must be possible to recombine them to produce a child with blue eyes and brown hair as the result of the cross.]

6. (Ergodicity) *It should be possible, through a finite sequence of applications of the genetic operators, to access any point in the search space* $\mathcal{C}$ *given any population* $\mathfrak{B}(t)$.
This provides the *raison d'être* for the mutation operator.
[Even if the whole population has blue eyes, it must be possible to produce a brown-eyed child. The mutation operator usually ensures this.]

5 Random, Respectful Recombination

Given a set Ξ of formae, an obvious question is whether it is possible to construct a recombination operator which simultaneously respects and properly assorts the formae, and if so, whether there is more than one such operator. It is simple to show that not all sets of formae can be so respected and properly assorted; a set which cannot be is described in section 9. Those which can be are said to be *separable*, and a recombination operator which respects and properly assorts a set of formae is said to *separate* them.

The principle of respect amounts loosely to the requirement that characteristics shared by both parents are passed on to their children. It is useful to define the *similarity set* $\eta \oplus \zeta$ of chromosomes η and ζ as the highest precision forma which contains them both:

$$\eta \oplus \zeta \equiv \bigcap \left\{ \xi \in \Xi \;\middle|\; \eta, \zeta \in \xi \right\} .$$

In the familiar case of schemata, this is the schema having the alleles which η and ζ share at its (only) definition points. For example,

$$\begin{array}{r} 1\,0\;1\;1\;1\;0\;0\;1 \\ \oplus\; 1\,0\;0\;1\;0\;0\;1\;0 \\ \hline 1\,0\;\square\;1\;\square\;0\;\square\;\square \end{array}$$

Respect then amounts to the requirement that all children produced by recombining η and ζ be members of their similarity set:

$$\forall \eta, \zeta \in \mathcal{C}\ \forall a \in \mathcal{A}_X : X(\eta, \zeta, a) \in \eta \oplus \zeta .$$

Clearly if Ξ is separable (that is, capable of simultaneous respect and proper assortment) the children required for proper assortment must also lie in the similarity set of the parents. It follows that a recombination operator which, given any pair of parents η and ζ, returns a randomly-selected member of $\eta \oplus \zeta$, is guaranteed both to respect and properly assort the formae. The recombination operator which makes a *uniform* random choice of children from the similarity set of the parents is called *random, respectful recombination*, which we abbreviate to R^3 (figure 3).

For example, when crossing the binary chromosomes 1010 and 0011, each of the four members of the similarity set $\square 01 \square$ (0010, 0011, 1010, and 1011) is chosen with probability one quarter by R^3.

6 Gene Transmission

It is instructive to notice that in the case of schemata for binary string representations the R^3 operator reduces to uniform crossover (Radcliffe [9]); see Syswerda [11] and Eshelman *al* [3] for details of uniform crossover) but that this is not so for k-ary representations with $k > 2$. For while uniform crossover requires each allele to be selected

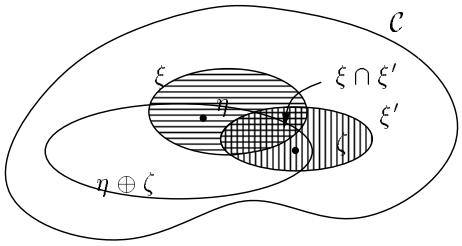

Figure 3: Any pair of chromosomes η and ζ have a *similarity set* denoted $\eta \oplus \zeta$, which is the smallest forma from Ξ containing them both. Respect requires that any children produced by recombining two solutions lie in their similarity set. The R^3 operator makes a uniform random choice of child from the parents' similarity set. If the set of formae, Ξ, is separable (i.e. capable of simultaneous respect and proper assortment) then R^3 will separate them.

randomly from one of the parents, R^3, after copying all the shared alleles to the child, fills in the rest of the chromosome with genes randomly selected from the allele sets.

In constructing the *cycle* crossover operator for permutation representations and *o*-schemata, the principal aim of Oliver *et al* [8] was to ensure that every allele in the child was taken from one or other of the parents. More recently, Whitley *et al* [12] have used a similar criterion applied to *edges* rather than vertices of the graph to construct a highly-successful *edge recombination* operator for the TSP. It would be natural, therefore, to attempt to formulate a seventh design principle which specifies that all alleles present in the child are to be transmitted from one or other of its parents.

The difficulty with this for general representations and formae is that the notion of an allele is not necessarily well-defined (e.g. *locality formae*, introduced in section 10). We can, however, make some progress by introducing the notion of a *complete orthogonal basis* for a set of equivalence relations Ψ. In the case of schemata and *o*-schemata the basis E we shall seek to construct consists of all the relations with a single definition point. In the case of four-gene schemata, this gives

$$E = \{\ \blacksquare\square\square\square,\ \square\blacksquare\square\square,\ \square\square\blacksquare\square,\ \square\square\square\blacksquare\ \}.$$

We can then define the intersection of compatible relations in the obvious way (figure 4) so that

$$\blacksquare\square\square\cdots\square \cap \square\blacksquare\square\cdots\square = \blacksquare\blacksquare\square\cdots\square.$$

This then allows us to propose the seventh design principle as follows:

7. (Strict Transmission) *Given a complete orthogonal basis $E \subset \Psi$ for a set of equivalence relations Ψ over the search*

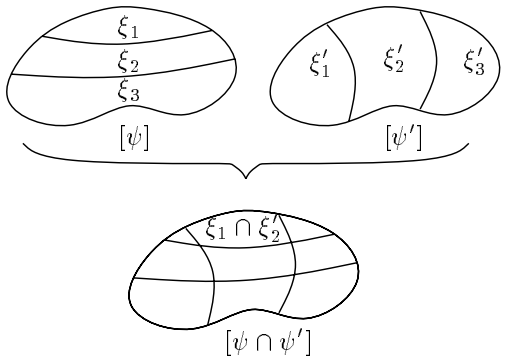

Figure 4: The set of formae induced by the equivalence relations ψ, ψ' and $\psi \cap \psi'$. The formae induced by $\psi \cap \psi'$ are intersections of those induced by ψ and ψ'.

space, under each equivalence relation in the basis E, every child produced by recombination must be equivalent to one of its parents.
A recombination operator which obeys this principle will be said to be *strictly transmitting.* In the familiar case, this is precisely the requirement that every allele in the child come from one parent or the other.
[If one parent has blue eyes and the other has brown eyes, the child must have blue or brown eyes.]

Formalising the notion of a basis is quite hard, and the hurried reader can safely skip the rest of this section if uninterested in the formalism.

We begin by defining intersection for two equivalence relations. For these purposes an equivalence relation $\sim$ is best described by a binary function

$$\psi : \mathcal{C} \times \mathcal{C} \longrightarrow \{\, 0, 1 \,\}$$

which returns 1 if its arguments are equivalent and 0 if they are not:

$$\psi(\eta, \zeta) = \begin{cases} 1, & \text{if } \eta \sim \zeta, \\ 0, & \text{otherwise.} \end{cases}$$

We can then define the intersection of two equivalence relations $\psi, \psi' \in \Psi$ by

$$(\psi \cap \psi')(\eta, \zeta) = \begin{cases} 1, & \text{if } \psi(\eta, \zeta) = \psi'(\eta, \zeta) = 1, \\ 0, & \text{otherwise.} \end{cases}$$

Given this, a subset $E \subset \Psi$ will be said to be a complete orthogonal basis for Ψ provided that

- (Completeness) All relations $\psi \in \Psi$ can be constructed as as intersection of the basic relations:

$$\forall \psi \in \Psi \; \exists \, E_\psi \subset E : \bigcap E_\psi = \psi .$$

- (Orthogonality) Every forma F induced by every basic relation $\psi \in E$ is compatible with every forma F' induced by every other basic relation $\psi' \in E$:

$$\forall \psi, \psi' \in E \;\; (\psi \neq \psi') \\ \forall F \in [\psi] \; \forall F' \in [\psi'] : F \cap F' \neq \emptyset \, ,$$

where $[\psi]$ is the set of equivalence classes (formae) induced by ψ.

(binary schemata)	respect	properly assort	strictly transmit
1-point	•		•
2-point	•		•
1-pt shuffle	•	•	•
2-pt shuffle	•	•	•
uniform/R^3	•	•	•

Table 1: Operators for binary schemata

(k-ary schemata)	respect	properly assort	strictly transmit
1-point	•		•
2-point	•		•
1-pt shuffle	•	•	•
2-pt shuffle	•	•	•
uniform	•	•	•
R^3	•	•	

Table 2: Operators for k-ary schemata

The transmission principle described above is now well-defined for a general set of formae provided that the equivalence relations inducing these formae have a complete orthogonal basis associated with them. We can define a *gene* as a basic equivalence relation and an allele as one of the equivalence classes (basic formae) induced by such a basic equivalence relation, motivating the term *gene transmission.*

7 Schemata

It is illuminating to classify the standard genetic operators used for recombining k-ary chromosomes and schemata. Tables 1 & 2 show which of the operators respect, properly assort, and strictly transmit schemata for binary and k-ary chromosomes respectively.

Notice that every standard crossover operator respects schemata and is strict in gene transmission, but that traditional 1- and $2-$point crossover do not properly assort schemata. Plainly, just as R^3 for binary chromosomes is identical to uniform crossover, if one modifies it by enforc-

ing strict gene transmission then uniform crossover is recovered for higher k-ary chromosomes also. These points are noteworthy principally because of the contrasting situation for (admittedly more complex) permutation representations and *o*-schemata described in the next section.

It should also be pointed out that while proper assortment and strict transmission have been defined as properties which recombination operators either possess (with respect to a given set of formae) or do not, in reality there are degrees of assortment and transmission. We might, for example, say that traditional crossover (with any number of cross-points) *weakly assorts* schemata on the basis that given a finite number of crosses between the parents and their various intermediate children it is possible to generate a child having any admixture of the parents' genes. Similarly, operators which do not enforce strict transmission of genes from parents to children will nevertheless pass on genes with some finite probability, giving rise to the notion of *partial transmission.*

8 O-schemata

In their paper on the travelling sales-rep problem, Goldberg & Lingle [5] introduced both the partially-mapped crossover (PMX) operator for permutations and the original *o*-schemata. Subsequently Oliver *et al* [8] introduced a variation of *o*-schemata and Goldberg [4] introduced several others. For present purposes we shall consider only the original *o*-schemata, which are most similar to schemata; similar analysis is possible for the other types.

Goldberg's *o*-schemata are induced by equivalence relations which relate chromosomes having the same elements of the permutation in particular positions on the chromosome. This seems like a useful and appropriate relation when the absolute positions of the objects labelled by the permutation matter. In job-shop scheduling, for example, the numbers typically represent jobs and the positions on the chromosome specify on which machines and in which order the jobs should be placed. In these cases the absolute positions seem relevant and one might reasonably expect operators which reliably respect, assort and transmit *o*-schemata to be of great use. As Whitley *et al* [12] have argued, however, it is not apparent that absolute city position is of any great significance in the TSP; this is discussed further in the next section.

Table 3 shows the way in which *o*-schemata are manipulated by four standard crossover operators for permutations and the R^3 operator. The operators are Goldberg's PMX crossover [5], Oliver *et al*'s cycle crossover [8], Davis's order crossover operator, modified as described in [8], and what Davis calls *uniform (permutation) crossover,* which relates to order crossover in exactly the same way as uniform crossover relates to traditional 1- and 2-point crossover. (The elements from one parent are copied wherever the binary mask that acts as the control parameter has a 1, and the remaining elements are used to fill the gaps in the order that they occur in the other parent.)

o-schemata	respect	properly assort	strictly transmit
PMX	•		
cycle	•		•
order			
uniform			
R^3	•	•	

Table 3: Operators for *o*-schemata

The R^3 operator in this case acts simply by inserting the common genes straight into the child chromosome and then arranging the remaining elements of the permutation at random in the gaps.

9 Edge Formae

It seems clear, as Whitley *et al* [12] have argued, that the edges rather than the vertices of the graph are central to the TSP. While there might be some argument as to whether or not the edges should be taken to be directed, the symmetry of the euclidean metric used in the evaluation function suggests that undirected edges suffice.

If the towns (vertices) in an n-city TSP are numbered 1 to n, and the edges are described as non-ordered pairs of vertices (a, b), then apparently suitable *edge formae* are simply sets of edges, subject to the condition that no vertex appears in the description of more than two edges. Unfortunately, these formae are not separable. To see this, consider two tours η and ζ, with η containing the fragment 2–1–3 and ζ containing 4–1–3. Plainly these have the common edge $(1, 3)$ $[\equiv (3, 1)]$. We shall describe the formae by listing the edges they require to be present in angle-brackets, so that η is an instance of the forma $\langle(1, 2)\rangle$ and ζ is an instance of the forma $\langle(1, 4)\rangle$. These formae are clearly compatible, because any tour containing the fragment 2–1–4 is in their intersection[4]

$$\langle(1, 2)\rangle \cap \langle(1, 4)\rangle = \langle(1, 2), (1, 4)\rangle.$$

However, any recombination operator which respected the formae would be bound to include the common edge $(1, 3)$ in all offspring from these parents, thus precluding generating a child in $\langle(1, 2), (1, 4)\rangle$. Since proper assortment

[4] Curiously, the intersection operation for these edge formae looks like the set *union* operation. This is because $\langle(1, 3)\rangle$ is really an abbreviation for the set of chromosomes containing the 1–3 edge.

requires that this child be capable of being generated this shows that these formae are not separable.

When Whitley *et al* introduced their powerful and attractive *edge recombination* operator they argued that 'there is no need for any new notion of "schema", with its own special schema theorem' because edge recombination manipulates an 'underlying binary representation' in the usual way. They went on to assert that 'where the parents have the same [edge], the offspring will have the same edge'. From the description in the paper, however, given tours containing the fragments shown above, it appears that while edge recombination *would* always generate a legal tour, there is no constraint which *requires* it always to transmit the common edge to the child. The operator does, however, provide a high rate of transmission of edges, this having been the major design criterion.

We can, of course, define the R^3 operator for the edge formae, even though they are not separable: it works simply by copying common edges into the child and then putting in random edges in such a way as to complete a legal tour. The lack of separability simply ensures that R^3 does not properly assort the formae.

10 Locality Formae

All of the formae discussed thus far have been fairly similar to traditional schemata. We now introduce *locality formae*, (Radcliffe [10]) which are rather different in character. Locality formae relate chromosomes on the basis of their closeness to each other. Suppose our function is defined over a real interval $[a, b)$. We then define formae which divide the interval up into strips of arbitrary width. Thus, a forma is a half-open interval $[\alpha, \beta)$ with α and β both lying in the range $[0, b - a]$. These formae are separable. Respect requires that all children are instances of any formae which contain both parents η and ζ. Clearly the similarity set of η and ζ (the smallest interval which contains them both) is $[\eta, \zeta]$, where we have assumed, without loss of generality, that $\zeta \geq \eta$. Thus respect requires that all their children lie in $[\eta, \zeta]$. Similarly, if η is in some interval $\xi = [\alpha, \beta)$ and ζ lies in some other interval $\xi' = [\alpha', \beta')$, then for these formae to be compatible the intersection of the intervals that define them must be non-empty ($\beta > \alpha'$; figure 5) and so picking a random element from the similarity set $[\eta, \zeta]$ allows an element to be picked which lies in the intersection, showing that R^3 fulfils the requirements of proper assortment (figure 6).

The n-dimensional R^3 operator picks a random point in the n-dimensional hypercuboid with corners at the two chromosomes η and ζ (figure 7). This operator has been tested on De Jong's functions [2] which are not all obviously suitable for locality formae, and performed surprisingly well,

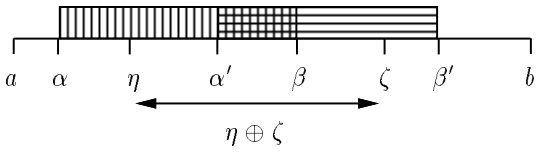

Figure 5: Given $\eta \in [\alpha, \beta)$ and $\zeta \in [\alpha', \beta')$, with $\zeta > \eta$, the formae are compatible only if $\beta > \alpha'$. The arrow shows the similarity set $\eta \oplus \zeta$.

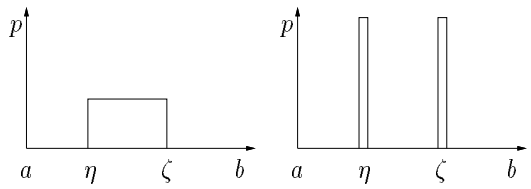

Figure 6: The left-hand graph shows (schematically) the probability of selecting each point along the axis under R^3 ("top hat"). The right-hand graph shows the corresponding diagram for standard crossover with real genes.

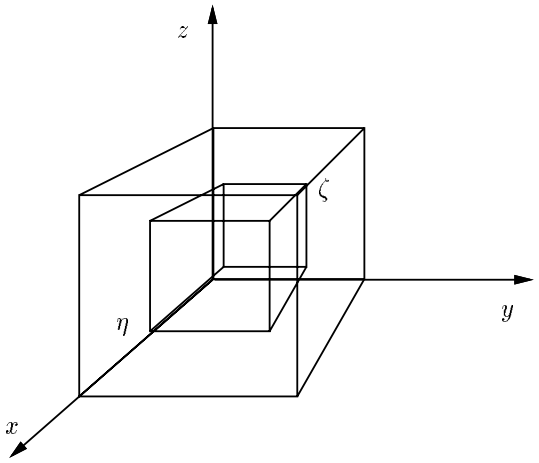

Figure 7: The n-dimensional R^3 operator for real genes picks any point in the hypercuboid with corners at the chromosomes being recombined, η and ζ.

apparently out-performing standard binary representations on four of the five functions. Full results are given in [9] and [10].

Both this operator and its natural analogue for k-ary string representations, which for each locus picks a random value in the range defined by the alleles from the two parents, suffer from a bias away from the ends of the interval. It is therefore necessary to introduce a mutation operator which

offsets this bias in order to satisfy the ergodicity condition expressed in principle 6. An appropriate mutation operator acts with very low probability to introduce the extremal values at an arbitrary locus along the chromosome. In the one dimensional case this amounts to occasionally replacing the value of one of the chromosomes with an a or a b. The combination of R^3 and such *end-point* mutation appears to provide a surprisingly powerful set of genetic operators for some problems.

Locality formae are not, of course, the only alternatives to schemata which can be applied to real-valued problems, and there is no suggestion here that locality formae should be seen as a generic or definitive alternative to schemata. It would be interesting, for example, to attempt to construct formae and representations on the basis of fourier analysis, or some other complete orthonormal set of functions over the space being searched.

11 Future Directions

The random, respectful recombination operator discussed above has been introduced as one which automatically respects and properly assorts separable formae. In some circumstances R^3 is useful in its own right, but its principal utility seems likely to lie in providing a starting point from which to construct more sophisticated operators. The most obvious way to do this is to modify the flat probability distribution over the similarity set of the parents which R^3 uses. Such modification could either exclude some children entirely or simply bias the search towards some subset of the similarity set on the basis of other information or intuitions about the structure of the search space.

The problems which motivated the ideas in this paper are ones for which the author was unable to find traditional schemata which characterised the search spaces in useful ways—graph optimisation, neural networks and the travelling sales-rep problem among others. Little emphasis has been placed on gene transmission in this paper, even though this appears at first a very natural constraint to place on recombination operators. This lack of emphasis derives from a lack of clarity as to the meaning of the term "gene" in the context of problems such as those listed. The construction of a complete orthogonal basis for a set of formae provides a mechanism for defining genes rigorously, after which normal schema theory may be applied. There are, however, sets of formae for which no orthogonal basis exists. It is in areas such as these that the ideas in this paper are most likely be useful, through the development of equivalence relations (and their associated formae) which well characterise the regularities in the search space. This may be helpful even when these formae do not admit the construction of a complete orthogonal basis, and so do not allow a subsequent return to schema analysis.

12 Conclusion

Formae have been shown to be useful generalisations of schemata which help the exploitation of intrinsic parallelism in non string-based problems and extend the scope of the "fundamental" (schema) theorem. The random, respectful recombination operator (R^3) has been introduced as an operator which is sometimes useful in its own right and might often be a useful starting point for developing more sophisticated operators for a range of problems.

Acknowledgements

I would like to thank Mike Norman for many useful discussions about genetic algorithms, and also Andrew J. S. Wilson and Mike for careful criticism of a draft of this paper.

Some of this work was supported by the Edinburgh Parallel Computing Centre, a multidisciplinary project which receives major grants from the Department of Trade and Industry, the Computer Board and the Science and Engineering Research Council. The author acknowledges support from the University of Edinburgh and from Industrial Affiliates.

References

[1] J. David Schaffer (ed), *Proceedings of the Third International Conference on Genetic Algorithms,* Morgan Kaufmann (San Mateo) 1989.

[2] Kenneth A. De Jong, *A Genetic-Based Global Function Optimization Technique,* Technical Report 80-2, University of Pittsburgh, 1980.

[3] Larry J. Eshelman, Richard A. Caruna, & J. David Schaffer, *Biases in the Crossover Landscape,* in [1].

[4] D. E. Goldberg, *Genetic Algorithms in Search, Optimization & Machine Learning,* Addison-Wesley (Reading, Mass) 1989.

[5] D. E. Goldberg, & Robert Lingle Jr, *Alleles, Loci and the Travelling Salesman Problem,* in *Proceedings of an International Conference on Genetic Algorithms,* Lawrence Erlbaum Associates (Hillsdale) 1985.

[6] John Grefenstette, Rajeev Gopal, Brian Rosmaita, & Dirk Van Gucht, *Genetic Algorithms for the Travelling Salesman Problem,* in *Proceedings of an International Conference on Genetic Algorithms,* Lawrence Erlbaum Associates (Hillsdale) 1985.

[7] J. H. Holland, *Adaptation in Natural and Artificial Systems,* University of Michigan Press (Ann Arbor) 1975.

[8] I. M. Oliver, D. J. Smith & J. R. C. Holland, *A Study of Permutation Crossover Operators in the Traveling*

Salesman Problem, in *Proceedings of the Second International Conference on Genetic Algorithms,* Lawrence Erlbaum Associates (Hillsdale, NJ) 1987.

[9] N. J. Radcliffe, *Genetic Neural Networks on MIMD Machines,* Ph.D. Thesis, Edinburgh University 1990.

[10] N. J. Radcliffe, *Equivalence Class Analysis of Genetic Algorithms,* to appear in *Complex Systems.*

[11] Gilbert Syswerda, *Uniform Crossover in Genetic Algorithms,* in [1].

[12] Darrell Whitley, Timothy Starkweather & D'Ann Fuquay, *Sheduling Problems and Traveling Salesmen: The Genetic Edge Recombination Operator* in [1].