SIMPLE GENETIC ALGORITHMS WITH LINEAR FITNESS

Michael D. Vose

Computer Science Dept., 107 Ayres Hall

The University of Tennessee

Knoxville, TN 37996-1301

vose@cs.utk.edu

Alden H. Wright[1]

Computer Science Dept.

The University of Montana

Missoula, MT. 59812

wright@cs.umt.edu

ABSTRACT

A general form of stochastic search is described (*random heuristic search*) and some of its general properties are proved. This provides a framework in which the simple genetic algorithm (SGA) is a special case. The framework is used to illuminate relationships between seemingly different probabilistic perspectives of SGA behavior. Next, the SGA is formalized as an instance of random heuristic search. The formalization is then used to show expected population fitness is a Lyapunov function in the infinite population model when mutation is zero and fitness is linear. In particular,

[1]This work was done while visiting the computer science department of the University of Tennessee.

the infinite population algorithm must converge, and average population fitness increases from one generation to the next. The consequence for a finite population SGA is that the expected population fitness increases from one generation to the next. Moreover, the only stable fixed point of the expected next population operator corresponds to the population consisting entirely of the optimal string. This result is then extended by way of a perturbation argument to allow nonzero mutation.

1 Introduction

The classical genetic algorithm[2], referred to as the *simple genetic algorithm (SGA)*, is stochastic. It cannot be characterized as converging to any particular population, nor is it possible to predict population trajectories or to know the change in population fitness from one generation to the next. It is natural, therefore, to approach the SGA from a probabilistic perspective.

At one extreme are complete Markov models (see, for example, Davis 1991; Nix and Vose 1991; Vose 1995) which capture every aspect of SGA behavior. This type of formalization is complex however, and extracting desired information can be nontrivial.

A seemingly more tractable approach is to study the SGA's *expected transition* from one generation to the next. This is given by a deterministic function $\mathcal{G}$ which from the current population p produces the expected next generation $\mathcal{G}(p)$.

It will be proved in this paper that within the framework of *random heuristic search*, the function $\mathcal{G}$ is independent of population size. It therefore simultaneously describes (in a sense that will be made precise) the expected next generation for all population sizes. A further consequence of the framework

[2] A fixed length binary GA, with crossover, mutation, and proportional selection.

is that $\mathcal{G}(p)$ is also the sampling distribution according to which population p should be sampled to produce the SGA's *actual* (as opposed to expected) next generation. It follows that everything that could ever be proved about a finite population SGA corresponds to some property of $\mathcal{G}$.

Another approach is to make the simplifying assumption of an infinite population. This is equivalent to the assumption that the SGA will, at every generation, move to the expected next generation. Consequently, an infinite population is not prerequisite to drawing inferences about SGA behavior on the basis of this model. Population trajectories in the infinite population case are determined by following the expected transition

$$p,\ \mathcal{G}(p),\ \mathcal{G}^2(p),\ \ldots$$

(see Juliany and Vose 1994 for some interesting partial visualizations of this dynamical system). This approach is similar to the infinite population models of population genetics (see, for example, Nagylaki, 1992). Because of the various interpretations of $\mathcal{G}$ pointed out above, this dynamical system also encodes complete information about SGA behavior.

This paper makes these various interpretations of $\mathcal{G}$ clear by way of the general framework for stochastic search referred to above as random heuristic search. The SGA will be formalized as an instance of random heuristic search, and the formalization will be used to obtain the main results of this paper based on the infinite population model.

For the class of *linear fitness functions* (defined in a later section), it is shown that the average population fitness increases from one generation to the next under the influence of selection and crossover (mutation is not included). Population trajectories in the infinite population model converge to a fixed point of $\mathcal{G}$, and all fixed points correspond to populations consisting entirely of a single string type. The consequence for a finite population SGA is that the expected population fitness increases from one generation to the next. Moreover, the only stable fixed point of the expected next population operator corresponds to the population consisting entirely of the optimal string. This result is then generalized to include nonzero mutation.

Rabinovich and Wigderson (1991) prove a stronger result under the additional restrictions that populations must be invariant under permutations of the bits of their members, there is no mutation, and the fitness function is counting-ones. They give a constructive proof that average fitness converges to optimum fitness in logarithmic time. Related to this is a result of Vose (1994) that under suitable conditions (which are satisfied in the case of linear fitness) the infinite population model of any instance of random heuristic search – in particular the SGA – converges in logarithmic time.

2 Notation

The set of integers is denoted by Z, and the set of integers modulo 2 is denoted by Z_2. The symbol $\Re$ denotes the set of real numbers. For any collection C of real numbers or real valued functions, αC denotes the collection whose members are those of C multiplied by α.

Angle brackets $<\cdots>$ denote a tuple which is to be regarded as a column vector. The column vector of all 1s is denoted by $\mathbf{1}$, and the zero vector is denoted by 0. The $n \times n$ identity matrix is I_n, and the jth column of I_n is the vector e_j. For vector x, $diag(x)$ denotes the square diagonal matrix with iith entry x_i. Indexing of vectors and matrices begins with 0.

The symbol "$\cdot$" is used as a place holder, except when occurring between column vectors as in $x \cdot y$ where it indicates the componentwise product of x and y (i.e., $diag(x)y$). Transpose is indicated with superscript T. The standard vector norm is $\|x\| = \sqrt{x^T x}$. The open ball of radius ε about the element or set x is denoted by $\mathcal{B}_\varepsilon(x)$, the closed ball by $\overline{\mathcal{B}}_\varepsilon(x)$.

Composition of functions f and g is $f \circ g(x) = f(g(x))$. The ith iterate f^i of f is defined by

$$\begin{aligned} f^0(x) &= x \\ f^{i+1}(x) &= f \circ f^i(x) \end{aligned}$$

Modulus (or absolute value) is denoted by $|\cdot|$. The notation $O(f)$ denotes a function (with similar domain and codomain as f), call it g, such that $|\, g\,| \leq c\,|\, f\,|$ for some constant c. The notation $o(f)$ represents a function (with similar domain and codomain as f), call it h, such that $|\,h\,|\,/\,|\,f\,| \to 0$. In the case where f is a vector or matrix, $|\cdot|$ is to be interpreted as a norm.

Curly brackets $\{\cdots\}$ are used as grouping symbols and to specify both sets and multisets. Square brackets $[\cdots]$ are, besides their standard use as specifying a closed interval of real numbers, used to denote an *indicator function*: if *expr* is an expression which may be true or false, then

$$[expr] = \begin{cases} 1 & \text{if } expr \text{ is true} \\ 0 & \text{otherwise} \end{cases}$$

The delta function is $\delta_{i,j} = [i = j]$. The equivalence of objects x and y is indicated by $x \equiv y$.

3 Random Heuristic Search

3.1 Representation

Random heuristic search can be thought of as an initial collection of elements P_0 chosen from some search space Ω of cardinality n together with some *transition rule* τ which from P_i will produce another collection P_{i+1}. In general, τ will be iterated to produce a sequence of collections

$$P_0 \xrightarrow{\tau} P_1 \xrightarrow{\tau} P_2 \xrightarrow{\tau} \ldots$$

The beginning collection P_0 is referred to as the *initial population*, the first population (or *generation*) is P_1, the second generation is P_2, and so on. Populations are multisets.

Not all transition rules are allowed. Obtaining a good representation for populations is a first step towards characterizing admissible τ. Define the *simplex* to be the set

$$\Lambda \;=\; \{< x_0, ..., x_{n-1} > : \mathbf{1}^T x = 1,\ x_j \geq 0\}$$

An element p of Λ corresponds to a population according to the following rule for defining its components

$$p_j \;=\; \text{the proportion in the population of the } j\text{th element of } \Omega$$

For example, if $\Omega = \{0, 1, 2, 3, 4, 5\}$ then $n = 6$. The population $\{1, 0, 3, 1, 1, 3, 2, 2, 4, 0\}$ would be represented by the vector $p = < .2, .3, .2, .2, .1, .0 >$ given table 1.

coordinate	corresponding element of Ω	percentage of P_0
p_0	0	2/10
p_1	1	3/10
p_2	2	2/10
p_3	3	2/10
p_4	4	1/10
p_5	5	0/10

Table 1

The cardinality of each generation $P_0, P_1, \ldots$ is a parameter r called the *population size*. Hence the proportional representation given by p unambiguously determines a population once r is known. The vector p is referred to as a *population vector*. The distinction between population and population vector will often be blurred, because the population size is fixed. In particular, τ may be thought of as mapping the current population vector to the next.

To get a feel for the geometry of the representation space, Λ is shown in figure 1 for $n = 2, 3, 4$. The diagrams represent Λ (a line segment, triangle, and solid tetrahedron). The arrows show the coordinate axes of the ambient space (the projection of the coordinate axes are being viewed in the second diagram, which is three dimensional, and in the last diagram where the ambient space is four dimensional).

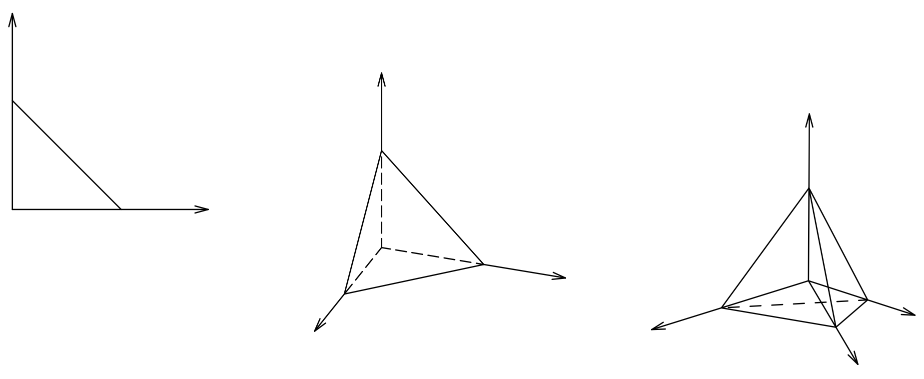

Figure 1

In general, Λ is a tetrahedron of dimension $n-1$ contained in an ambient space of dimension n. Note that each vertex of Λ corresponds to a unit basis vector of the ambient space; Λ is their convex hull. For example, the vertices of the solid tetrahedron (the right most diagram in figure 1) are at the basis vectors $e_0 =< 1, 0, 0, 0 >$, $e_1 =< 0, 1, 0, 0 >$, $e_2 =< 0, 0, 1, 0 >$, and $e_3 =< 0, 0, 0, 1 >$. Assuming that $\Omega = \{0, 1, 2, 3\}$, they correspond (respectively) to the following populations: r copies of 0, r copies of 1, r copies of 2, and r copies of 3. The center diagram will later be used as a schematic for general Λ, representing it for arbitrary n.

It should be realized that not every point of Λ corresponds to a finite population. In fact, only those rational points with common denominator r correspond to populations of size r. They are

$$\frac{1}{r} X_n^r \quad = \quad \frac{1}{r} \{< x_0, \ldots, x_{n-1} >: \ x_j \in Z, \ x_j \geq 0, \ \mathbf{1}^T x = r\}$$

For example, the points corresponding to $\frac{1}{4} X_4^4$ ($n = 4$ and $r = 4$) are the dots in figure 2.

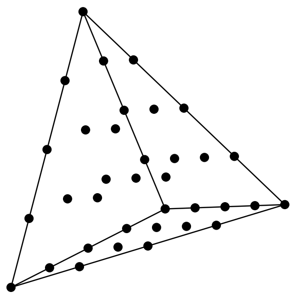

Figure 2

As $r \to \infty$, these rational points become dense in Λ. Since a rational point may represent arbitrarily large populations, a point p of Λ carries little information concerning population size. A natural view is therefore that Λ corresponds to populations of indeterminate size. This is but one of several useful interpretations. Another is that Λ corresponds to sampling distributions over Ω: since the components of p are nonnegative and sum to 1, p may be viewed as indicating that i is sampled with probability p_i.

In summary, random heuristic search appears to be a *discrete dynamical system* on Λ through the identification of populations with population vectors. That is, there is some transition rule

$$\tau : \Lambda \longrightarrow \Lambda$$

and what is of interest is the sequence of iterates beginning from some initial population vector p

$$p, \quad \tau(p), \quad \tau^2(p), \quad \ldots$$

This view is incomplete however, because the transitions are in general nondeterministic and not all transition rules are allowed. Next the stochastic nature of τ will be explained and admissible τ will be characterized.

3.2 Nondeterminism

Given the current population vector p, the next population vector $\tau(p)$ cannot be predicted with certainty because τ is stochastic. It is most conveniently thought of as resulting from r independent, identically distributed random choices. Let $\mathcal{G} : \Lambda \to \Lambda$ be a *heuristic function* (heuristic for short) which given the current population p produces a vector whose ith component is the probability that the ith element of Ω is chosen (with replacement). That is, $\mathcal{G}(p)$ is that probability vector which specifies the sampling distribution by which the aggregate of r choices forms the next generation. A transition rule τ is admissible if it corresponds to a heuristic function $\mathcal{G}$ in this way. Figure 3 depicts the relationship between p, Λ, Ω, $\mathcal{G}$, and τ through a sequence of generations (the illustration does not correspond literally to any particular case, it depicts how transitions between generations take place in general).

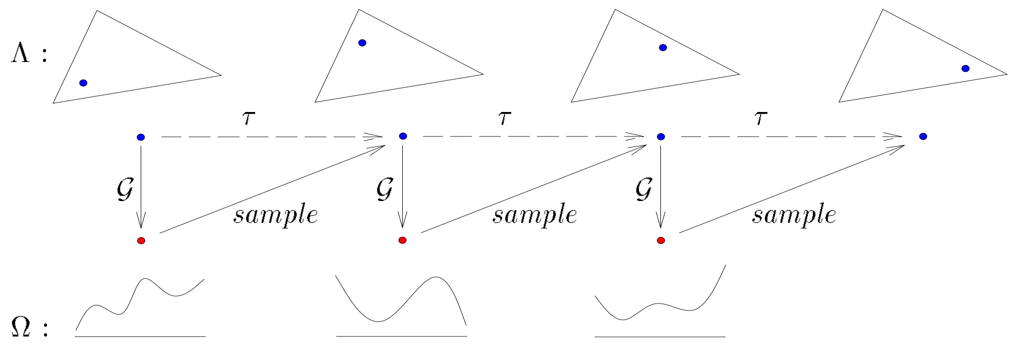

Figure 3

The triangles along the top row represent Λ, one for each of four generations. Each Λ contains a dot representing a population. These same populations are also represented in the second row with dots; τ maps from one to the next. The transition arrow for τ is dashed to indicate that it is an induced map, computed by following the solid arrows. The third row of dots are images of populations under $\mathcal{G}$. Below each is a curve, suggesting the sampling distribution over Ω which it represents. The line segments in the bottom row represent Ω.

The transition from one generation to the next proceeds a follows. First $\mathcal{G}$ is applied to produce a vector which represents a sampling distribution (curve) over Ω. Next, r independent samples, with replacement, are made from Ω according to this distribution (represented in figure 3 by *"sample"*) to produce the next generation.

For example, let $\Omega = \{0, 1, 2, 3\}$ and suppose the heuristic is

$$\mathcal{G}(p) \quad = \quad < 0, p_1, 2p_2, 3p_3 > / \sum i p_i$$

Let the initial population be $p = < .25, .25, .25, .25 >$. Since $\mathcal{G}(p) = < 0, 1/6, 1/3, 1/2 >$, the probability of sampling 0 is 0, of sampling 1 is 1/6, of sampling 2 is 1/3, and of sampling 3 is 1/2. With population size $r = 100$, the transition rule corresponds to making 100 independent samples, with replacement, according to these probabilities.

A plausible next generation is therefore $\tau(p) = <0, .17, .33, .50>$. Note that the sampling distribution $\mathcal{G}(p)$ used in forming the next generation $\tau(p)$ depends on the current population p. Going one generation further, the new current population is $\tau(p)$ and the sampling distribution for producing the next generation is given by $\mathcal{G}(\tau(p)) \approx < 0, .07296, .28326, .64377>$. It is therefore plausible that the second generation is $\tau^2(p) = <0, .07, .28, .65>$.

Note the conceptually dual interpretation of Λ. It serves as both the space of populations and as the space of probability distributions over Ω. The first natural question is: What is the expected next generation? The following identity, referred to as the *multinomial theorem* will be useful in answering this question:

$$(\mathbf{1}^T x)^k = k! \sum_{v \in X_n^k} \prod_{j<n} \frac{x_j^{v_j}}{v_j!}$$

Theorem 3.1 *Let p be the current population vector. The expected next population vector is $\mathcal{G}(p)$.*

Sketch of proof: Without loss of generality, $\Omega = \{0, \ldots, n-1\}$. The first step is to determine for each possible population vector the probability that it represents the next generation. Feasible populations are $\frac{1}{r} X_n^r$. To obtain a general population vector $q = v/r$, it must happen that v_0 choices out of r are 0, which has probability

$$\begin{pmatrix} r \\ v_0 \end{pmatrix} (\mathcal{G}(p)_0)^{v_0}$$

and v_1 choices out of the remaining $r - v_0$ must be 1, which has probability

$$\begin{pmatrix} r - v_0 \\ v_1 \end{pmatrix} (\mathcal{G}(p)_1)^{v_1}$$

and so on until finally v_{n-1} choices out of the remaining $r - v_0 - \cdots - v_{n-2}$ must be $n-1$, which has probability

$$\binom{r - v_0 - \cdots - v_{n-2}}{v_{n-1}} (\mathcal{G}(p)_{n-1})^{v_{n-1}}$$

The product of these probabilities reduces (after expanding the binomial coefficients) to

$$r! \prod_{j<n} \frac{(\mathcal{G}(p)_j)^{v_j}}{v_j!}$$

It follows that the expectation is given by

$$r! \sum_{v \,\in\, X_n^r} \frac{v}{r} \prod_{j<n} \frac{(\mathcal{G}(p)_j)^{v_j}}{v_j!}$$

Applying the operator $\sum e_i x_i \frac{\partial}{\partial x_i}$ to both sides of the multinomial theorem yields

$$k\, x\, (\mathbf{1}^T x)^{k-1} \;=\; k! \sum_i e_i\, x_i \sum_{v \in X_n^k} \frac{v_i x_i^{v_i - 1}}{v_i!} \prod_{j \neq i} \frac{x_j^{v_j}}{v_j!} \;=\; k! \sum_{v \in X_n^k} v \prod_j \frac{x_j^{v_j}}{v_j!}$$

Using this formula to simplify the expectation completes the proof. □

Observe that the statement of theorem 3.1 is independent of r. It therefore holds independent of population size. According to the law of large numbers, if the next generation's population vector q were obtained as the result of an infinite sample from the distribution described by $\mathcal{G}(p)$, then q would match the expectation, hence $q = \mathcal{G}(p)$. Because this corresponds to random heuristic search with an infinite population, the algorithm resulting from "$\tau = \mathcal{G}$"[3] is called the *infinite population algorithm.*

At this point, random heuristic search and the infinite population algorithm have been defined. The following sections will show how the heuristic $\mathcal{G}$ may be instantiated to yield the simple genetic algorithm.

To summarize:

[3]Strictly speaking, τ produces the next generation from the current, while $\mathcal{G}$ would in this context produce the *representation* of the next generation from the *representation* of the current. This distinction is conveniently blurred.

- By the observations made above, the dynamical system p, $\mathcal{G}(p)$, $\mathcal{G}^2(p)$, $\ldots$ represents population trajectories in the infinite population case.
- By construction, $\mathcal{G}(p)$ is the sampling distribution according to which population p should be sampled to produce the next generation in the finite population case.
- By theorem 3.1, $\mathcal{G}(p)$ is the expected next generation.
- Because theorem 3.1 is independent of r, it follows that $\mathcal{G}$ simultaneously describes the expected next generation for all population sizes.

4 The Simple Genetic Algorithm

Naturally enough, the Simple Genetic Algorithm is realized by specifying the search space Ω and identifying the heuristic function $\mathcal{G}$. Given the formalism provided by the previous section, nothing else is required to define the SGA; it is merely a special case of random heuristic search. Before proceeding, an algebraic framework will be established which will be used extensively throughout the paper.

4.1 Algebra

For positive integer ℓ, the set of length ℓ binary strings is the Cartesian product

$$\Omega = \underbrace{Z_2 \times \cdots \times Z_2}_{\ell \text{ times}}$$

Since the ℓ-digit binary representations of integers in the interval $[0, 2^\ell)$ coincide with the elements of Ω, they are regarded as being the same. To make explicit the dependence of Ω on ℓ, it may be written as ${}^m\Omega$. When elements of Ω are written as strings, the standard practice of putting least significant bit to the right is followed. For example, if $\ell = 3$ then $n = 2^\ell$ and

$$\Omega \;=\; \{000, 001, 010, 011, 100, 101, 110, 111\} \;\equiv\; \{0, 1, 2, 3, 4, 5, 6, 7\}$$

Elements of Z_2 form a finite field under the operations of addition and multiplication modulo 2

$$\begin{array}{c|cc} \oplus & 0 & 1 \\ \hline 0 & 0 & 1 \\ 1 & 1 & 0 \end{array} \qquad\qquad \begin{array}{c|cc} \otimes & 0 & 1 \\ \hline 0 & 0 & 0 \\ 1 & 0 & 1 \end{array}$$

These operations are extended to Ω by applying them coordinate-wise. By convention, $\otimes$ takes precedence over $\oplus$, and both bind more tightly than operations which are not modulo 2. When representing elements of Ω as column vectors, the least significant bit is at the top. This unusual ordering pays large dividends in the form of simplified formulas and theorems. Continuing in the context of the example above,

$$3 \oplus 5 \;\equiv\; 011 \oplus 101 \;\equiv\; <1 \oplus 1, 1 \oplus 0, 0 \oplus 1> \;\equiv\; 110 \;\equiv\; 6$$

$$6 \otimes 2 \;\equiv\; 110 \otimes 010 \;\equiv\; <0 \otimes 0, 1 \otimes 1, 1 \otimes 0> \;\equiv\; 010 \;\equiv\; 2$$

For $x \in \Omega$, let $\overline{x}$ abbreviate $\mathbf{1} \oplus x$. In standard computer science nomenclature, $\oplus$ is *exclusive-or* on integers, $\otimes$ is *and*, and $x \mapsto \overline{x}$ is *not*. Note that $\otimes$ distributes over $\oplus$. Because $\oplus$ differs from inclusive-or, De Morgan's laws do not hold with respect to $\{^-, \oplus, \otimes\}$.

The set $\{x \in \Omega : x \otimes k = k\}$ is denoted by Ω_k. Each $i \in \Omega$ has a unique representation $i = u \oplus v$ where $u \in \Omega_k$ and $v \in \Omega_{\overline{k}}$. This follows from the identity

$$i \;=\; i \otimes k \,\oplus\, i \otimes \overline{k} \;=\; u \oplus v$$

Let σ_k be the permutation matrix defined by

$$(\sigma_k)_{i,j} \quad = \quad [i \oplus j = k]$$

The permutation σ_k corresponds to applying the map $i \mapsto i \oplus k$ to subscripts. That is,

$$\sigma_k < x_0, \ldots, x_{n-1} > \quad = \quad < x_{0 \oplus k}, \ldots, x_{(n-1) \oplus k} >$$

It is easily verified that the σ_k are symmetric and $\sigma_i \sigma_j = \sigma_{i \oplus j}$. It follows that the σ_k commute and are self inverse.

4.2 Selection

The symbol s will be used for three equivalent (though different) things. This overloading of s does not take long to get used to because context makes meaning clear. The benefits are clean and elegant presentation and the ability to use a common symbol for ideas whose differences are often conveniently blurred.

First, $s \in \Lambda$ can be regarded as a *selection distribution* describing the probability s_i with which i is selected (with replacement) from the current population for participation in forming the next generation. A selected element is an intermediate step towards producing the next population, not typically a member of it. In total, $2r$ such selections will be made, the aggregate of which is sometimes referred to as the *gene pool.*

Second, $s : \Lambda \to \Omega$ can be regarded as a *selection function* which is nondeterministic. The result $s(p)$ of applying s to p is i with probability given by the ith component s_i of the selection distribution. Of course, for there to be a nontrivial dependence on p, the selection distribution must be some function $\mathcal{F}$ of p. The function $\mathcal{F}$ is referred to as the *selection scheme.*

Third, $s \in \Lambda$ can be regarded as a population vector.

In analogy with survival of the fittest, an integral part of $\mathcal{F}$ is a *fitness function* $f : \Omega \to \Re$ which can be used (in a variety of ways) to determine a selection scheme. The fitness function is assumed to be injective. The value $f(i)$ is called the *fitness* of i. Through the identification $f_i = f(i)$, the fitness function may be regarded as a vector.

Proportional selection refers to the selection function corresponding to the selection scheme

$$\mathcal{F}(x) \quad = \quad f{\cdot}x / f^T x$$

(recall $f{\cdot}x$ denotes $diag(f)x$). When proportional selection is being used, it is assumed that the fitness function is positive.

By letting the heuristic $\mathcal{G}$ be the selection scheme, results from previous sections apply to selection. For example, with population size $2r$, $\tau(p)$ becomes the gene pool. Invoking theorem 3.1, the expected gene pool is described by the population vector $s = \mathcal{F}(p)$. By definition, the selection distribution is $s = \mathcal{F}(p)$. Hence, as elements of Λ, the selection distribution is identical to the expected gene pool.

4.3 Mutation

The symbol μ will also be used for three different (though related) things.

First, $\mu \in \Lambda$ can be regarded as a distribution describing the probability μ_i with which i is selected to be a *mutation mask* (additional details will follow).

Second, $\mu : \Omega \to \Omega$ can be regarded as a *mutation function* which is nondeterministic. The result

$\mu(x)$ of applying μ to x is $x \oplus i$ with probability given by the ith component μ_i of the distribution μ. The i occurring in $x \oplus i$ is referred to as a mutation mask. The application of μ to x is referred to as *mutating* x.

Third, $\mu \in [0, 0.5)$ can be regarded as a *mutation rate* which implicitly specifies the distribution μ according to the rule

$$\mu_i = (\mu)^{\mathbf{1}^T i} (1-\mu)^{\ell - \mathbf{1}^T i}$$

The distribution μ need not correspond to any mutation rate, although that is certainly the classical situation. Any element $\mu \in \Lambda$ whatsoever is allowed.

The effect of mutating x using mutation mask i is to alter the bits of x in those positions where the mutation mask i is 1. When mutation is affected by a rate, the probability of selecting mask i depends only on the number of 1s that i contains.

If the mutation rate is nonzero (the typical case), then every element of Ω has a positive probability of being the result of $\mu(x)$. Mutation is said to be *zero* if $\mu_i = \delta_{i,0}$. For arbitrary $\mu \in \Lambda$, mutation is called *positive* if $\mu_i > 0$ for all i.

Mutation is called *independent* if for all j and k

$$\mu_j = \sum_{k \otimes i = 0} \mu_{i \oplus j} \sum_{\overline{k} \otimes i = 0} \mu_{i \oplus j}$$

Independent mutation is of interest because of its relationship to the crossover operator.

Theorem 4.1 *If μ corresponds to a mutation rate, then μ is independent.*

Sketch of proof: Replacing the dummy variable i by $i \oplus j$, the rightmost factor in the definition of independence becomes

$$\sum_{\overline{k}\otimes(i\oplus j)=0} \mu_i \;=\; (1-\mu)^\ell \sum_{\overline{k}\otimes i=\overline{k}\otimes j} \left(\frac{\mu}{1-\mu}\right)^{\mathbf{1}^T i}$$

Note that if $i = \overline{k} \otimes j \oplus k \otimes u$ then $\overline{k} \otimes i = \overline{k} \otimes j$. Hence summation over the range of i appropriate to the right hand side above corresponds to summation over $u \in \Omega_k$. Letting $m = \mathbf{1}^T k$, the right hand side is

$$(1-\mu)^\ell \sum_{u\,\in\,\Omega_k} \left(\frac{\mu}{1-\mu}\right)^{\mathbf{1}^T\overline{k}\otimes j+\mathbf{1}^T u} \;=\; (1-\mu)^{\ell-m} \left(\frac{\mu}{1-\mu}\right)^{\mathbf{1}^T\overline{k}\otimes j} \sum_{u\,\in\,{}^m\Omega} \mu^{\mathbf{1}^T u}(1-\mu)^{m-\mathbf{1}^T u}$$

Since the rightmost sum above is $(\mu + (1-\mu))^m = 1$ (by the binomial theorem), it follows that

$$\sum_{k\otimes i=0} \mu_{i\oplus j} \sum_{\overline{k}\otimes i=0} \mu_{i\oplus j} \;=\; (1-\mu)^{\ell-\mathbf{1}^T k} \left(\frac{\mu}{1-\mu}\right)^{\mathbf{1}^T\overline{k}\otimes j} (1-\mu)^{\ell-\mathbf{1}^T\overline{k}} \left(\frac{\mu}{1-\mu}\right)^{\mathbf{1}^T k\otimes j} \;=\; \mu_j$$

$\square$

4.4 Crossover

It is convenient to use the concept of *partial probability* . Let $\zeta : A \to B$ and suppose $\phi : A \to [0,1]$. To say "$\xi = \zeta(a)$ with partial probability $\phi(a)$" means that $\xi = b$ with probability $\sum_a [\zeta(a) = b]\, \phi(a)$.

The description of crossover parallels the description of mutation given in the previous section; the symbol χ will be used for three different (though related) things.

First, $\chi \in \Lambda$ can be regarded as a distribution describing the probability χ_i with which i is selected to be a *crossover mask* (additional details will follow).

Second, $\chi : \Omega \times \Omega \to \Omega$ can be regarded as a *crossover function* which is nondeterministic. The result $\chi(x, y)$ is $x \otimes i \oplus \overline{\imath} \otimes y$ with partial probability $\chi_i/2$ and is $y \otimes i \oplus \overline{\imath} \otimes x$ with partial probability $\chi_i/2$. The i occurring in the definition of $\chi(x, y)$ is referred to as a crossover mask. The application of $\chi(x, y)$ to x, y is referred to as *recombining* x and y.

The arguments x and y of the crossover function are called *parents*, the pair $x \otimes i \oplus \overline{i} \otimes y$ and $y \otimes i \oplus \overline{i} \otimes x$ are referred to as their *children*. Note that crossover produces children by exchanging the bits of parents in those positions where the crossover mask i is 1. The result $\chi(x, y)$ is called their *child*.

Third, $\chi \in [0, 1]$ can be regarded as a *crossover rate* which specifies the distribution χ according to the rule

$$\chi_i = \begin{cases} \chi c_i & \text{if } i > 0 \\ 1 - \chi + \chi c_0 & \text{if } i = 0 \end{cases}$$

where the distribution $c \in \Lambda$ is referred to as the *crossover type*. Classical crossover types include *1-point crossover*, for which

$$c_i = \begin{cases} 1/(\ell - 1) & \text{if } \exists k \in (0, \ell) \;.\; i = 2^k - 1 \\ 0 & \text{otherwise} \end{cases}$$

and *uniform crossover*, for which $c_i = 2^{-\ell}$. However, any element $c \in \Lambda$ whatsoever is allowed.

The following theorem says that when mutation is independent, it may be performed either before or after crossover; the effects are exactly the same.

Theorem 4.2 *If μ is independent, then for arbitrary χ*

$$prob\{\chi(\mu(x), \mu(y)) = z\} = prob\{\mu(\chi(x, y)) = z\}$$

Sketch of proof: The left hand probability is obtained by summing over all combinations of mutation masks and crossover masks which yield $\chi(\mu(x), \mu(y)) = z$,

$$\sum_{i,j,k} \mu_i \, \mu_j \, \chi_k \, \frac{1}{2}([(x \oplus i) \otimes k \oplus \overline{k} \otimes (y \oplus j) = z] + [(x \oplus i) \otimes \overline{k} \oplus k \otimes (y \oplus j) = z])$$

Splitting the sum across the "+", performing the change of variable $k \mapsto \overline{k}$ and recombining yields

$$\sum_{i,j,k} \mu_i \, \mu_j \, \frac{\chi_k + \chi_{\overline{k}}}{2} \, [(x \oplus i) \otimes k \oplus \overline{k} \otimes (y \oplus j) = z]$$

Since $(x \oplus i) \otimes k \oplus \overline{k} \otimes (y \oplus j) = z$ is equivalent to $x \otimes k \oplus \overline{k} \otimes y = z \oplus h$ where $h = i \otimes k \oplus j \otimes \overline{k}$, organizing the sum according to values of h yields

$$\sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} \sum_h [x \otimes k \oplus \overline{k} \otimes y = z \oplus h] \sum_{i \otimes k \oplus j \otimes \overline{k} = h} \mu_i \mu_j$$

Note that $i \otimes k \oplus j \otimes \overline{k} = h$ is equivalent to the pair of relations $k \otimes i = k \otimes h$ and $\overline{k} \otimes j = \overline{k} \otimes h$. Thus the last sum above is

$$\sum_{k \otimes (i \oplus h) = 0} \mu_i \sum_{\overline{k} \otimes (j \oplus h) = 0} \mu_j$$

which, by the independence of μ, is μ_h. Substituting μ_h for this sum yields

$$\sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} \sum_h [x \otimes k \oplus \overline{k} \otimes y = z \oplus h] \, \mu_h$$

Splitting the sum across the "+", changing variables, and recombining as before yields

$$\sum_{k,h} \mu_h \, \chi_k \, \frac{1}{2}([x \otimes k \oplus \overline{k} \otimes y = z \oplus h] + [x \otimes \overline{k} \oplus k \otimes y = z \oplus h])$$

which is equal to $\Pr\{\mu(\chi(x, y)) = z\}$. □

Obtaining child z from parents x and y via the process of mutation and crossover is called *mixing* and has probability denoted by $m_{x,y}(z)$. Formulas for these probabilities, contained in the preceding proof, are recorded in the following theorem.

Theorem 4.3 *If mutation is performed before crossover, then*

$$m_{x,y}(z) \quad = \quad \sum_{i,j,k} \mu_i \, \mu_j \, \frac{\chi_k + \chi_{\overline{k}}}{2} \, [(x \oplus i) \otimes k \oplus \overline{k} \otimes (y \oplus j) = z]$$

If mutation is performed after crossover, then

$$m_{x,y}(z) \quad = \quad \sum_{j,k} \mu_j \, \frac{\chi_k + \chi_{\overline{k}}}{2} \, [x \otimes k \oplus \overline{k} \otimes y = z \oplus j]$$

Mixing has many symmetries. The most fundamental are:

Theorem 4.4 *Whether or not μ is independent, and whether or not mutation is performed before or after crossover,*

$$m_{x,y}(z) \quad = \quad m_{y,x}(z) \quad = \quad m_{x \oplus z, y \oplus z}(0)$$

Sketch of proof: Theorem 4.3 is symmetric in x and y. The indicator function in theorem 4.3 corresponding to mutation before crossover is equivalent to

$$[((x \oplus z) \oplus i) \otimes k \oplus \overline{k} \otimes ((y \oplus z) \oplus j) = 0]$$

The case of mutation after crossover is analogous. $\square$

The matrix M with i,jth entry $m_{i,j}(0)$ is called the *mixing matrix* . It is symmetric by theorem 4.4.

The *mixing scheme* $\mathcal{M} : \Lambda \to \Lambda$ is defined by the component equations

$$\mathcal{M}(x)_i \quad = \quad x^T \sigma_i M \sigma_i x$$

4.5 SGA's Heuristic

In procedural terms, the Simple Genetic Algorithm is defined as producing the next generation according to:

1. Obtain two parents by the selection function.
2. Mutate the parents by the mutation function.
3. Produce their child by the crossover function.
4. Put the result into the next generation.

5. If the next generation contains less than r members, go to step 1.

The corresponding heuristic $\mathcal{G}$ is the composition of mixing and selection

$$\mathcal{G} = \mathcal{M} \circ \mathcal{F}$$

For this to be correct, it suffices that i is chosen for the next generation with probability $\mathcal{G}(p)_i$, where p is the current population. The probability that i occurs at step 4 is

$$\begin{aligned}
&\Pr\{i = \chi(\mu(s(p)), \mu(s(p)))\} \;=\\
&\sum_{<x,y>} \Pr\{<x,y> = <s(p), s(p)>\} \Pr\{i = \chi(\mu(x), \mu(y))\} \;=\\
&\sum_{x,y} \mathcal{F}(p)_x \mathcal{F}(p)_y \sum_{u,v} \mu_{x\oplus u}\ \mu_{y \oplus v} \Pr\{i = \chi(u,v)\} \;=\\
&\sum_{x,y} \mathcal{F}(p)_x \mathcal{F}(p)_y \sum_{u,v} \mu_{x\oplus u}\ \mu_{y \oplus v} \sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} [u \otimes k \oplus \overline{k} \otimes v = i]
\end{aligned}$$

Making the change of variable $u \mapsto x \oplus u$ and $v \mapsto y \oplus v$ yields

$$\sum_{x,y} \mathcal{F}(p)_x \mathcal{F}(p)_y \sum_{u,v,k} \mu_u\, \mu_v \frac{\chi_k + \chi_{\overline{k}}}{2} [(x \oplus u) \otimes k \oplus \overline{k} \otimes (y \oplus v) = i]$$

Applying theorems 4.3 and 4.4 to the last sum gives

$$\sum_{x,y} \mathcal{F}(p)_x \mathcal{F}(p)_y\, M_{x \oplus i, y \oplus i}$$

which equals $(\mathcal{M} \circ \mathcal{F}(p))_i$ as required.

This completes the definition of the simple genetic algorithm, showing it to be an instance of random heuristic search and identifying its heuristic. The formalization presented here is a much more general and complete version of the original model developed in Vose 1990.

5 Linear Fitness

A fitness function is called *linear* if it has the form $f(i) = a^T i + b$ for some $a \in \Re^\ell$, $b \in \Re$. Since the SGA (as considered in this paper) uses proportional selection, is is assumed that f is injective and positive. This section uses the previous formalism to show average population fitness is a Lyapunov function when mutation is zero and fitness is linear. In particular, the infinite population algorithm corresponding to the simple genetic algorithm must converge to a vertex of Λ (i.e., to some e_j representing a population consisting entirely of string j) and average population fitness increases from one generation to the next. The consequence for a finite population SGA is that the expected population fitness increases from one generation to the next. Moreover, the only stable fixed point of the expected next population operator corresponds to the population consisting entirely of the optimal string.

Random heuristic search algorithms are classified according to the behavior of their heuristic functions. An instance of random heuristic search is *focused* if $\mathcal{G}$ is continuously differentiable and for every $p \in \Lambda$ the sequence

$$p, \quad \mathcal{G}(p), \quad \mathcal{G}(\mathcal{G}(p)), \quad \ldots$$

converges. In this case, $\mathcal{G}$ is also said to be focused. In terms of search, this condition implies that path determined by following what the heuristic is expected to produce will lead to some state x. By the continuity of $\mathcal{G}$,

$$\mathcal{G}(x) \quad = \quad \mathcal{G}(\lim_{l \to \infty} \mathcal{G}^{(l)}(p)) \quad = \quad \lim_{l \to \infty} \mathcal{G}^{(l+1)}(p) \quad = \quad x$$

Hence such points x satisfy $\mathcal{G}(x) = x$ and are called *fixed points* of $\mathcal{G}$. Note that, since the SGAs heuristic is infinitely differentiable, $\mathcal{G}$ in this case is focused provided its iterates converge.

There do not seem to be many techniques for establishing that iterates of a function must converge. When there is a single fixed point, the contraction mapping theorem (see Loomis & Sternberg, 1968) is a natural candidate. But in the case of several fixed points, that method is not as easily applied. The following discrete analogue of Lyapunov's stability theorem is often useful in situations where the contraction mapping theorem is not.

Lemma 5.1 *If $\mathcal{G}$ has finitely many fixed points and ϕ is a continuous function satisfying*

$$x \neq \mathcal{G}(x) \quad \Longrightarrow \quad \phi(x) \ > \ \phi(\mathcal{G}(x))$$

then $\mathcal{G}$ is focused.

Sketch of proof: Let U_j be an $\varepsilon > 0$ ball about fixed point ω_j, with ε sufficiently small that $x \in U_j \ \Rightarrow \ \mathcal{G}(x) \notin \ U_k$ for $k \neq j$. The proof proceeds by contradiction; thus the sequence $\mathcal{G}(x)$, $\mathcal{G}^{(2)}(x)$, $\mathcal{G}^{(3)}(x)$, $\ldots$ lies in the complement of $\bigcup U_j$ infinitely often. By compactness, let z be a limit point in this complement. Hence $\phi(z) \ > \ \phi(\mathcal{G}(z))$. By continuity, let V be a closed neighborhood of z such that $\phi(z) \ > \ \phi(\mathcal{G}(V))$. Since z is the limit of a sequence of points in $\bigcup_k \mathcal{G}^{(k)}(V)$, this leads to the contradiction $\phi(z) \ > \ \phi(z)$. $\square$

The function ϕ occurring above is called a *Lyapunov function.* The condition on ϕ given in the proposition may be taken as $x \neq \mathcal{G}(x) \implies \phi(x) < \phi(\mathcal{G}(x))$ since it is actually the monotone behavior of ϕ along trajectories that matters.

In the special case of linear fitness and zero mutation, $\phi(x) = f^T x$ is a Lyapunov function. Note that $\phi(x)$ is the average population fitness, since

$$f^T x \ = \ \sum f(i)\, x_i \ = \ \sum (\text{fitness of } i)(\text{proportion of } i \text{ in population } x)$$

Theorem 5.2 *If fitness is linear and mutation is zero, then $\mathcal{G}$ is focused, the only fixed points are vertices of Λ, and average population fitness increases from one generation to the next.*

Before proving this theorem, some preliminary results will be established.

Proposition 5.3 *For $a \in \Re^\ell$ and $i \in \Omega$,*

$$a^T i \;=\; \frac{1}{2}\sum a_j (1 - (-1)^{i \otimes 2^j})$$

Sketch of proof: Using the identity $x = \frac{1}{2}(1-(-1)^x)$, valid for $x \in \{0,1\}$, yields

$$a^T i \;=\; \sum a_j\, i_j \;=\; \sum a_j\, i^T 2^j \;=\; \frac{1}{2}\sum a_j (1 - (-1)^{i^T 2^j})$$

□

Proposition 5.4

$$\sum_h 2M_{u\oplus h, v\oplus h} - \delta_{u,h} - \delta_{v,h} = 0$$

Sketch of proof: By the definition of M, the sum is

$$\begin{aligned}
& \sum_h \sum_{i,j,k} \mu_i \mu_j \frac{\chi_k + \chi_{\overline{k}}}{2} [0 = (u \oplus h \oplus i) \otimes k \oplus \overline{k} \otimes (v \oplus h \oplus j)] - \delta_{u,h} - \delta_{v,h} \\
= \;& -2 + \sum_{i,j,k} \mu_i \mu_j \frac{\chi_k + \chi_{\overline{k}}}{2} \sum_h [0 = h \otimes k \oplus h \otimes \overline{k} \oplus (u \oplus i) \otimes k \oplus \overline{k} \otimes (v \oplus j)] \\
= \;& -2 + \sum_{i,j,k} \mu_i \mu_j \frac{\chi_k + \chi_{\overline{k}}}{2} \sum_h [h = (u \oplus i) \otimes k \oplus \overline{k} \otimes (v \oplus j)] \\
= \;& -2 + \sum_{i,j,k} \mu_i \mu_j \frac{\chi_k + \chi_{\overline{k}}}{2} \\
= \;& -2 + 2
\end{aligned}$$

□

Proposition 5.5 *If mutation is zero, then*

$$\sum_i (-1)^{i^T 2^j} M_{u\oplus i, v\oplus i} \;=\; (-1)^{u^T 2^j} [u \otimes 2^j = v \otimes 2^j]$$

Sketch of proof: Replacing the dummy summation variable i by $i \oplus u$, and using the fact that the summation is symmetric in u and v shows the left hand side is

$$(-1)^{u^T 2^j} \sum_i (-1)^{i^T 2^j} M_{i,u\oplus v\oplus i} \quad = \quad (-1)^{v^T 2^j} \sum_i (-1)^{i^T 2^j} M_{i,u\oplus v\oplus i}$$

Since this equation cannot hold when $u^T 2^j \neq v^T 2^j$ unless both sides are zero, it follows that the sum is

$$\begin{aligned}
& (-1)^{u^T 2^j} [u \otimes 2^j = v \otimes 2^j] \sum_i (-1)^{i^T 2^j} M_{i,u\oplus v\oplus i} \\
= \; & (-1)^{u^T 2^j} [u \otimes 2^j = v \otimes 2^j] \sum_i (-1)^{i^T 2^j} \sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} [0 = i \otimes k \oplus \overline{k} \otimes (u \oplus v \oplus i)] \\
= \; & (-1)^{u^T 2^j} [u \otimes 2^j = v \otimes 2^j] \sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} \sum_i (-1)^{i^T 2^j} [i = \overline{k} \otimes (u \oplus v)] \\
= \; & (-1)^{u^T 2^j} \sum_k \frac{\chi_k + \chi_{\overline{k}}}{2} \left([u \otimes 2^j = v \otimes 2^j] \, (-1)^{(\overline{k} \otimes (u \oplus v))^T 2^j} \right) \\
= \; & (-1)^{u^T 2^j} [u \otimes 2^j = v \otimes 2^j] \sum_k \frac{\chi_k + \chi_{\overline{k}}}{2}
\end{aligned}$$

□

The proof of theorem 5.2 will now be sketched. Let $q = \mathcal{G}(p)$ and note that

$$(f^T p)^2 \, q_i \quad = \quad (f^T p) \, f_i p_i \; + \; \sum_{u,v} f_u p_u f_v p_v (M_{u\oplus i, v\oplus i} - \delta_{v,i})$$

Multiplying by f_i and summing over i yields

$$(f^T p)^2 \, f^T q \quad = \quad (f^T p) \sum_i f_i^2 p_i \; + \; \sum_{u,v} f_u p_u f_v p_v \sum_i f_i (M_{u\oplus i, v\oplus i} - \delta_{v,i})$$

Using the inequality $\sum f_i^2 p_i \geq \sum^2 f_i p_i$ and dividing through by $(f^T p)^2$ gives

$$f^T q \quad \geq \quad f^T p \; + \; s^T A s$$

where $s = \mathcal{F}(p)$ and $A_{u,v} = \sum f_i \, (M_{u\oplus i, v\oplus i} - \delta_{v,i})$. Hence $f^T x$ is a Lyapunov function if $s^T A s \geq 0$ and if $\sum f_i^2 p_i = \sum^2 f_i p_i$ only at fixed points. It will turn out that $s^T A s = 0$.

Note that $\sum f_i^2 p_i - \sum^2 f_i p_i$ is the variance of a random variable which takes value f_i with probability p_i. It therefore can be zero only if $p_i = \delta_{i,k}$ for some k (fitness is injective). Hence fixed points are vertices of Λ (i.e., the e_j representing populations consisting entirely of string j) since elsewhere this variance is positive and $f^T q > f^T p$.

Since $2\,s^T A s = s^T(A^T + A)s$, the matrix A may be replaced by $A^T + A$. The proof is completed by showing $A = 0$. Using propositions 5.3, 5.4, and 5.5,

$$\begin{aligned}
A_{u,v} &= \sum_i (b + a^T i)\,(2M_{u\oplus i, v\oplus i} - \delta_{u,i} - \delta_{v,i}) \\
&= \sum_i \frac{1}{2} \sum_j a_j (1 - (-1)^{i^T 2^j})\,(2M_{u\oplus i, v\oplus i} - \delta_{u,i} - \delta_{v,i}) \\
&= -\frac{1}{2} \sum_i \sum_j a_j (-1)^{i^T 2^j}\,(2M_{u\oplus i, v\oplus i} - \delta_{u,i} - \delta_{v,i}) \\
&= -\frac{1}{2} \sum_j a_j \sum_i (-1)^{i^T 2^j}\,(2M_{u\oplus i, v\oplus i} - \delta_{u,i} - \delta_{v,i}) \\
&= -\frac{1}{2} \sum_j a_j \left(2\,(-1)^{u^T 2^j}[u \otimes 2^j = v \otimes 2^j] - (-1)^{u^T 2^j} - (-1)^{v^T 2^j}\right) \\
&= 0
\end{aligned}$$

6 The Small Mutation Case

If, at all fixed points $x \in \Lambda$ of $\mathcal{G}$, the differential $d\mathcal{G}_x$ of $\mathcal{G}$ at x has no eigenvalue of modulus 1, then $\mathcal{G}$ is called *hyperbolic*. This section generalizes previous results to include mutation by way of a perturbation argument. The idea is that a focused, hyperbolic heuristic will remain so if not too greatly perturbed. The proof makes use of the Lyapunov function ϕ (of the previous section) in an essential way.

This material is significantly more technical than what has come before. This is unavoidable; some things are by nature complicated. Basic background in calculus and topology is assumed (see, for example, Loomis & Sternberg, 1968; Akin, 1993). A simplified though complete picture of the central ideas of this section may be obtained by skipping over the propositions and the proofs of the lemmas.

Throughout it will be assumed that $\mathcal{G}$ is hyperbolic; hence $d\mathcal{G}_x$ is a hyperbolic linear map at fixed points x of $\mathcal{G}$. Assuming that an arbitrarily small perturbation of f is allowed, this assumption is justified – in the case of linear fitness and zero mutation – by the following theorem.[4]

Theorem 6.1 *Let $f(i) = a^T i + b$ be a positive and injective fitness function. If mutation is zero, then an arbitrarily small perturbation of b will preserve these properties of f while guaranteeing the hyperbolicity of $\mathcal{G}$.*

Sketch of proof: By theorem 5.2, fixed points of $\mathcal{G}$ are at vertices of Λ. The spectrum of $d\mathcal{G}_{e_k}$ is

$$\left\{ \frac{f_{i \oplus k}}{f_k} 2\, M_{0,i} \;:\; i = 1, 2, \ldots, n-1 \right\} \bigcup \{0\}$$

(see Vose and Wright, 1995). Note that the fractions f_i/f_j are all monotonic functions of b. Hence they take on an infinite number of distinct values as b varies in any small interval. It follows that $\mathcal{G}$ is hyperbolic for some perturbation of b since only finitely many values for each ratio can defeat hyperbolicity.[5] □

Let M refer to an arbitrary mixing matrix, and let M_0 refer to a mixing matrix corresponding to zero mutation. Because fixed points $x \in \Lambda$ and the eigenvalues of $d\mathcal{G}_x$ depend continuously on M, theorem 6.1 implies that $\mathcal{G}$ may be assumed to be hyperbolic for all $M \in \mathcal{B}_\lambda(M_0)$, provided that λ is sufficiently small.

[4] It is also justified in the general case where f need not be linear, but the proof is beyond the scope of this paper.

[5] Note that this same proof idea can be used to show that a may be perturbed instead of b. In fact, it may be concluded that those linear f for which $\mathcal{G}$ is hyperbolic form a dense open set.

Proposition 6.2 *Suppose $||M - M_0||$ is bounded, and consider $\mathcal{G}$ as depending on the parameter M and as having domain $\overline{\mathcal{B}}_\varepsilon(e_j)$. The function $o(y) = \mathcal{G}(x+y) - \mathcal{G}(x) - d\mathcal{G}_x(y)$ is Lipschitz in y, uniformly in x and M, with Lipschitz constant $o(1)$ as $\varepsilon \to 0$.*

Sketch of proof: Note that $d\mathcal{G}_x$ is a continuous function of x and M. By hypothesis, $||M - M_0|| \le k$, so x and M range over compact sets. Hence $d\mathcal{G}_x$ is uniformly continuous in x and M. Next,

$$\begin{aligned}
||o(y) - o(z)|| &= ||\int_0^1 \frac{d}{dt}\, o(ty + (1-t)z)\, dt\, || \\
&= ||\int_0^1 (d\mathcal{G}_{x+ty+(1-t)z} - d\mathcal{G}_x)(y - z)\, dt\, || \\
&\le ||y - z|| \int_0^1 ||d\mathcal{G}_{x+ty+(1-t)z} - d\mathcal{G}_x||\, dt \\
&\le ||y - z|| \sup_{0 \le t \le 1} ||d\mathcal{G}_{x+ty+(1-t)z} - d\mathcal{G}_x||
\end{aligned}$$

By the uniform continuity of $d\mathcal{G}_x$ in x and M, the supremum is arbitrarily small provided y and z are. $\square$

The sequence $x, \mathcal{G}(x), \mathcal{G}^2(x), \ldots$ is called the *orbit* of x. The following lemma concerning orbits is the cornerstone of the perturbation result.

Lemma 6.3 *There exists $\varepsilon > 0$ and $\lambda > 0$ such that for all $M \in \mathcal{B}_\lambda(M_0)$ and all j, if the orbit of y does not leave $\mathcal{B}_\epsilon(e_j)$ then the orbit of y converges.*

Sketch of proof: Let $T = d\mathcal{G}_{e_j}$ for $\mathcal{G}$ with mixing matrix M_0, and let, for mixing matrix M, $g(x) = \mathcal{G}(e_j) - Te_j + (d\mathcal{G}_{e_j} - T)(x - e_j) + o(x - e_j)$ where $o(x - e_j) = \mathcal{G}(x) - \mathcal{G}(e_j) - d\mathcal{G}_{e_j}(x - e_j)$. For the remainder of this proof, $\mathcal{G}$ denotes the heuristic with mixing matrix M.

Note that $\mathcal{G}(e_j) - Te_j$ is Lipschitz with constant 0 (since it is constant), $d\mathcal{G}_{e_j} - T$ is Lipschitz (since it is linear) with constant $||d\mathcal{G}_{e_j} - T|| \to 0$ as $M \to M_0$, and $o(x - e_j)$ is Lipschitz with constant

$o(1)$ for $x \in \mathcal{B}_\epsilon(e_j)$ as $\epsilon \to 0$ (proposition 6.2). Hence g is Lipschitz with arbitrarily small Lipschitz constant. Moreover,

$$\begin{aligned}
\mathcal{G}(x) &= \mathcal{G}(e_j) + d\mathcal{G}_{e_j}(x - e_j) + o(x - e_j) \\
&= T(x - e_j) + \mathcal{G}(e_j) + (d\mathcal{G}_{e_j} - T)(x - e_j) + o(x - e_j) \\
&= Tx + g(x)
\end{aligned}$$

Since T is a hyperbolic linear map, let E^+ and E^- be its stable and unstable subspaces (respectively). Let $\|\cdot\|$ be a norm adapted to T and let

$$\alpha \;=\; \max\{\, \|T\big|_{E^+}\|,\ \|(T\big|_{E^-})^{-1}\|\,\} \;<\; 1$$

Since all norms are equivalent, there exist $\varepsilon_j > 0$ and $\lambda_j > 0$ such that the Lipschitz constant of $g(x)$ is less than $(1-\alpha)/2$ for $x \in \mathcal{B}_{\varepsilon_j}(e_j)$ and $M \in \mathcal{B}_{\lambda_j}(M_0)$. Therefore, the composition of g with the retraction to $\mathcal{B}_{\varepsilon_j}(e_j)$ has lipschitz constant less than $1-\alpha$, so the stable manifold theorem applies to show that orbits (under $\mathcal{G} = T + g$) which are confined to $\mathcal{B}_{\varepsilon_j}(e_j)$ converge (the fixed points of $\mathcal{G}$ belong to $\mathcal{B}_{\varepsilon_j}(e_j)$ for sufficiently small λ). Now let $\varepsilon = \min \varepsilon_j$ and $\lambda = \min \lambda_j$. $\square$

Proposition 6.4 *Given $\eta > 0$ there exists $\lambda > 0$ and $\xi > 0$ such that for all $M \in \mathcal{B}_\lambda(M_0)$,*

$$x \in \Lambda \setminus \bigcup \mathcal{B}_\eta(e_j) \quad\Longrightarrow\quad \phi(\mathcal{G}(x)) > \phi(x) + \xi$$

Sketch of proof: First, the apparently weaker result corresponding to $\xi = 0$ is proved by way of contradiction. Define $g(x) = \phi(\mathcal{G}(x)) - \phi(x)$ and let $M_j \to M_0$ with corresponding $x_j \in \Lambda \setminus \bigcup \mathcal{B}_\eta(e_j)$ be such that $g(x_j) \le 0$. By compactness (and passing to a subsequence if necessary), assume that $x_j \to x \in \Lambda \setminus \bigcup \mathcal{B}_\eta(e_j)$. By continuity, $g(x) \le 0$ for $M = M_0$ which contradicts that ϕ is a Lyapunov function.

Hence $g(x, M)$ is positive on some compact set of the form

$$\Big(\Lambda \setminus \bigcup \mathcal{B}_\eta(e_j)\Big) \;\times\; \overline{\mathcal{B}}_\lambda(M_0)$$

Choosing positive ξ less than the minimum of g over this set finishes the proof.

$\square$

Lemma 6.5 *If ε is sufficiently small, then there exists $\eta > 0$ and $\lambda > 0$ such that for all $M \in \mathcal{B}_\lambda(M_0)$ and for all j the following holds: If the orbit of x enters $\mathcal{B}_\eta(e_j)$ and then leaves $\mathcal{B}_\varepsilon(e_j)$, say at y, then $\phi(y) > \phi(e_j) + \eta$.*

Sketch of proof: For all j, if $\eta' \to 0$ and $\lambda' \to 0$, it follows by continuity that $\mid \mathcal{G}(x) - e_j \mid \to 0$ for $x \in \mathcal{B}_{\eta'}(e_j)$ and $M \in \mathcal{B}_{\lambda'}(M_0)$. Thus let η' and λ' be sufficiently small so that for all j,

$$x \in \mathcal{B}_{\eta'}(e_j) \ \wedge \ M \in \mathcal{B}_{\lambda'}(M_0) \quad \Longrightarrow \quad \mathcal{G}(x) \in \mathcal{B}_\varepsilon(e_j)$$

Apply proposition 6.4 with $\eta = \eta'$ and let the resulting λ be λ''. Using the resulting ξ (from proposition 6.4), choose $\eta < \min\{\eta', \xi/2, \varepsilon\}$ and $\lambda''' < \min\{\lambda', \lambda''\}$ sufficiently small so that

$$x \in \mathcal{B}_\eta(e_j) \ \wedge \ M \in \mathcal{B}_{\lambda'''}(M_0) \quad \Longrightarrow \quad \phi(\mathcal{G}(x)) > \phi(e_j) - \xi/2$$

Finally, let $z \in \mathcal{B}_\varepsilon(e_j)$ be such that $\mathcal{G}(z) = y$, and apply proposition 6.4 once more, letting the resulting λ be λ''''. Choose $\lambda = \min\{\lambda''', \lambda''''\}$.

Note that $z \notin \mathcal{B}_{\eta'}(e_j)$ since $y \notin \mathcal{B}_\varepsilon(e_j)$. Hence if ε is sufficiently small, then

$$z \in \Lambda \setminus \bigcup \mathcal{B}_{\eta'}(e_j)$$

and therefore $\phi(y) = \phi(\mathcal{G}(z)) > \phi(z) + \xi$. Let w be the last point on the orbit of x contained in $\mathcal{B}_\eta(e_j)$ prior to z. Then $\phi(\mathcal{G}(w)) > \phi(e_j) - \xi/2$. Moreover, ϕ increases along the trajectory from $\mathcal{G}(w)$ to z since those points are contained in $\Lambda \setminus \bigcup \mathcal{B}_\eta(e_j)$. Therefore $\phi(z) > \phi(e_j) - \xi/2$ and hence $\phi(y) > \phi(e_j) + \xi/2 > \phi(e_j) + \eta$. $\square$

The next step is the definition of sets on which orbits are under control in the sense that lemma 6.5 can be usefully applied. Before proceeding, note that a consequence of lemma 6.5 is that if f_j is maximal (among f_i), then an orbit entering $\mathcal{B}_\eta(e_j)$ cannot leave $\mathcal{B}_\varepsilon(e_j)$. Otherwise, there would be some y such that $\phi(y) > f_j$, which is not possible.

Let ψ be the permutation such that $i < j \iff f(\psi_i) < f(\psi_j)$. For $\mathcal{G}$ with mixing matrix M_0, define, depending on parameters $\gamma_j > 0$,

$$\begin{aligned} U_j &= \{x : \lim_{k\to\infty} \mathcal{G}^k(x) = e_{\psi_j}\} \\ C_j &= \bigcup_{i\leq j} U_i \\ D_j &= C_j \setminus \bigcup_{i<j} \mathcal{B}_{\gamma_i}(D_i) \end{aligned}$$

Proposition 6.6 *The C_j are closed, and invariant under $\mathcal{G}$ with mixing matrix M_0. Moreover,*

$$D_j \;\subset\; U_j \;\subset\; C_j \;\subset\; \bigcup_{i\leq j} \mathcal{B}_{\gamma_i}(D_i)$$

Sketch of proof: Since the U_j are invariant under $\mathcal{G}$, the C_j are also. That the C_j are closed follows by reverse induction on j. The base case $j = n-1$ is trivial, since $C_{n-1} = \Lambda$. By the induction hypothesis, assume C_j is closed. Since it is invariant, the dynamical system corresponding to $\mathcal{G}$ can be restricted to C_j. Note that $\phi(e_{\psi_j})$ is maximal over vertices of Λ in C_j. Hence e_{ψ_j} is an attractor since iterates of $\mathcal{G}$ must converge to a vertex and ϕ increases along trajectories (theorem 5.2). Therefore the basin of attraction U_j of e_{ψ_j} is open in C_j, so its complement is closed. The complement is C_{j-1} since the U_i are disjoint. Moreover,

$$\begin{aligned} C_j &\subset (C_j \setminus \bigcup_{i<j} \mathcal{B}_{\gamma_i}(D_i)) \cup \bigcup_{i<j} \mathcal{B}_{\gamma_i}(D_i) \\ &= D_j \cup \bigcup_{i<j} \mathcal{B}_{\gamma_i}(D_i) \\ &\subset \bigcup_{i\leq j} \mathcal{B}_{\gamma_i}(D_i) \end{aligned}$$

Therefore

$$D_j \;=\; C_j \setminus \bigcup_{i<j} \mathcal{B}_{\gamma_i}(D_i) \;\subset\; C_j \setminus C_{j-1} \;=\; U_j \;\subset\; C_j$$

□

Proposition 6.7 *Let $\eta > 0$. There exists $\varepsilon > 0$ and $\lambda > 0$ such that for all $x \in \mathcal{B}_\varepsilon(D_j)$ and $M \in \mathcal{B}_\lambda(M_0)$, the orbit of x enters $\mathcal{B}_\eta(e_{\psi_j})$. Moreover, $\varepsilon > 0$ and $\lambda > 0$ are independent of D_i for $i > j$.*

Sketch of proof: Since D_j is contained in U_j (proposition 6.6), for each $x \in D_j$ there exists k_x such that $\|\mathcal{G}^{k_x}(x) - e_{\psi_j}\| < \eta$ for mixing matrix M_0. Since $\mathcal{G}(x)$ is continuous in x and M, there exist ε_x and λ_x such that if $y \in \mathcal{B}_{\varepsilon_x}(x)$ and $M \in \mathcal{B}_{\lambda_x}(M_0)$, then $\|\mathcal{G}^{k_x}(y) - e_{\psi_j}\| < \eta$ for mixing matrix M. The sets $\mathcal{B}_{\frac{1}{2}\varepsilon_x}(x)$ are an open cover of D_j. By compactness, let $\mathcal{B}_{\frac{1}{2}\varepsilon_{x_1}}(x_1), \ldots, \mathcal{B}_{\frac{1}{2}\varepsilon_{x_k}}(x_k)$ be a subcover (D_j is closed since it is the intersection of C_j with the complement of an open set). Define $\varepsilon = \min \frac{1}{2}\varepsilon_{x_i}$ and $\lambda = \min \lambda_{x_i}$.

Now if $x \in \mathcal{B}_\varepsilon(D_j)$, then $x \in \mathcal{B}_\varepsilon(y)$ for some $y \in D_j$ and y is contained in some element of the open cover, say $\mathcal{B}_{\frac{1}{2}\varepsilon_{x_i}}(x_i)$. Hence $x \in \mathcal{B}_{\varepsilon_{x_i}}(x_i)$, since $\|x - x_i\| \le \|x - y\| + \|y - x_i\| \le \varepsilon + \varepsilon_{x_i}$. By construction, $\|\mathcal{G}^{k_x}(x) - e_{\psi_j}\| < \eta$ for all $M \in \mathcal{B}_\lambda(M_0)$. □

Lemma 6.8 *Given $\eta > 0$ there exists $\lambda > 0$ such that for all $M \in \mathcal{B}_\lambda(M_0)$ and for all j, if $x \in \mathcal{B}_{\gamma_j}(D_j)$ then the orbit of x enters $\mathcal{B}_\eta(e_{\psi_j})$. Here the γ_j depend on η and may be taken sufficiently small that*

$$\sup_{x \in \mathcal{B}_{\gamma_j}(D_j)} \phi(x) \;<\; \phi(e_{\psi_j}) + \eta$$

Sketch of proof: Choose the parameters γ_j (defining the sets D_j) sequentially ($j = 0, 1, \ldots$) to depend on η as follows. Let γ_j be the ε from proposition 6.7, and let the corresponding λ be λ_j. Now let $\lambda = \min \lambda_j$. This determines the sets D_j up to the parameter η. Since the ε of proposition 6.7 can be arbitrarily small, the γ_j can be arbitrarily small as well.

It will next be shown, by way of contradiction, that the γ_j may be sufficiently small that

$$\sup_{y \in \mathcal{B}_{\gamma_j}(C_j)} \phi(y) \;<\; \phi(e_{\psi_j}) + \eta$$

This suffices to complete the proof since $D_j \subset C_j$. Let $\gamma_j \to 0$, let corresponding $y \in \mathcal{B}_{\gamma_j}(C_j)$ be such that $\phi(y) > \phi(e_{\psi_j}) + \eta/2$, and let $x \in C_j$ be a point minimally distant from y. By compactness (pass to a subsequence if necessary), assume $x_j \to x \in C_j$ as $\gamma_j \to 0$. By continuity of ϕ, it follows that $\phi(x) \ge \phi(e_{\psi_j}) + \eta/2$. This contradicts

$$\sup_{x \in C_j} \phi(x) \;=\; \sup_{i \le j} \sup_{x \in U_i} \phi(x) \;=\; \sup_{i \le j} \phi(e_{\psi_i}) \;=\; \phi(e_{\psi_j})$$

□

Finally, the perturbation result can be proved.

Theorem 6.9 *If the fitness function f is positive and linear, then, at the possible expense of an arbitrarily small perturbation to f (which may be chosen to preserve linearity), there exists $\lambda > 0$ such that $\mathcal{G}$ is focused for all $M \in \mathcal{B}_\lambda(M_0)$.*

Sketch of proof: Choose $\varepsilon' > 0$ so that lemma 6.3 applies, and let the corresponding λ be λ'. Choose $0 < \varepsilon \leq \varepsilon'$ so that lemma 6.5 applies, and let the corresponding λ be λ''. Using the resulting η (from lemma 6.5), apply lemma 6.8, and let the corresponding λ be λ'''. Now let $\lambda = \min\{\lambda', \lambda'', \lambda'''\}$.

The proof proceeds by showing that every orbit will eventually become trapped with in some $\mathcal{B}_\varepsilon(e_j)$, whereupon it converges (lemma 6.3).

Since $\Lambda = \bigcup U_j \subset \bigcup \mathcal{B}_{\gamma_j}(D_j)$ (proposition 6.6), every $x \in \Lambda$ is contained in some $\mathcal{B}_{\gamma_j}(D_j)$. Hence by lemma 6.8, the orbit of x enters $\mathcal{B}_\eta(e_{\psi_j})$. If the orbit does not leave $\mathcal{B}_\varepsilon(e_{\psi_j})$, then the proof is complete. Otherwise, lemma 6.5 implies that when it does leave, say at y, then $\phi(y) > \phi(e_{\psi_j}) + \eta$. By construction, $x \in \mathcal{B}_{\gamma_j}(D_j) \implies \phi(x) < \phi(e_{\psi_j}) + \eta$ (lemma 6.8). Therefore, $y \in \mathcal{B}_{\gamma_k}(D_k)$ for some $k > j$.

The previous paragraph may be summarized by: if an orbit, upon entering $\mathcal{B}_\eta(e_i)$, does not become trapped in $\mathcal{B}_\varepsilon(e_i)$, then it enters some $\mathcal{B}_\eta(e_k)$ where $f_k > f_i$. It follows that the orbit must eventually become trapped, since the alternative is an infinite increasing sequence $f_{i_1} < f_{i_2} < f_{i_3}, \ldots$ which is not possible. $\square$

To summarize these results and make related observations, the addition of small amounts of mutation has the following effects:

- The average population fitness still increases in the infinite population model from populations which are not too homogeneous (this is the message of proposition 6.4).

- Because of the alternate interpretations of $\mathcal{G}$, this means that, in the finite population case, the expected population fitness still increases from populations which are not too homogeneous.

- Population trajectories in the infinite population model still converge (this is the message of theorem 6.9).

- Since a small perturbation of the dynamical system will not alter dimensions of the stable/unstable spaces of the differentials at fixed points (by hyperbolicity and continuity), there is still a unique stable fixed point of the expected next population operator.

In fact, the proof technique establishes a more general result in that the fitness function can be simlutaneously perturbed with mutation; the linearity of fitness need not be maintained.

7 Conclusion

Random heuristic search (a general form of stochastic search) has been described and some of its general properties were proved. By showing the simple genetic algorithm to be a special case, relationships between seemingly different probabilistic perspectives on SGA behavior were explained and a rigourous formalization of the simple genetic algorithm was obtained.

The formalization was used to show expected population fitness is a Lyapunov function in the infinite population model when mutation is zero and fitness is linear. In particular, the infinite population algorithm must converge, and average population fitness increases from one generation to the next. The consequence for a finite population SGA is that the expected population fitness increases from one generation to the next. Moreover, the only stable fixed point of the expected next population operator corresponds to the population consisting entirely of the optimal string.

These results were extended, by way of a perturbation argument, to allow nonzero mutation. The result obtained is that given a positive hyberbolic linear fitness function and small mutation, the expected population fitness still increases from populations which are not too homogeneous, population trajectories in the infinite population model still converge, and there is a unique stable fixed point of the expected next population operator (which does not correspond to the population consisting entirely of the optimal string). In fact, the proof technique establishes a more general result in that the fitness function can be simlutaneously perturbed with mutation; the linearity of fitness need not be maintained.

8 Acknowledgements

This research was supported by the National Science Foundation: IRI-9224917.

9 References

Akin, Ethan. (1993). *The General Topology of Dynamical Systems*. American Mathematical Society.

Davis, T. (1991). Toward An Extrapolation Of The Simulated Annealing Convergence Theory Onto The Simple Genetic Algorithm. Ph.D. Thesis, University of Florida.

Juliany, J & Vose, M. D. (1994). The Genetic Algorithm Fractal. *Evolutionary Computation*, v. 2, no. 2, 165-180.

Loomis, L. H. & Sternberg, S. (1968). *Advanced Calculus.* Reading, Mass.: Addison – Wesley.

Nagylaki, T. (1992). *Introduction to Theoretical Population Genetics.* Berlin, Germany: Springer Verlag.

Nix, A. E. & Vose, M. D. (1992). Modeling Genetic Algorithms With Markov Chains. *Annals of Mathematics and Artificial Intelligence*, 5, 79-88.

Rabinovich & Wigderson. (1991). An Analysis of a Simple Genetic Algorithm. *Proceedings of the Fourth International Conference on Genetic Algorithms* 215-221. San Mateo, CA: Morgan Kaufmann.

Vose, M. D. (1990). Formalizing Genetic Algorithms. In *The Proceedings of Genetic Algorithms, Neural Nets, and Simulated Annealing Applied to Problems in Signal and Image Processing.* Glasgow, U.K.: Institute of Electrical and Electronics Engineers.

Vose, M. D. & Wright, A. (1994). Stability of Vertex Fixed Points and Applications. Submitted to *Evolutionary Computation.*

Vose, M. D. (1994). Logarithmic Convergence of Random Heuristic Search. Submitted to *Evolutionary Computation.*

Vose, M. D. (in press). *The Simple Genetic Algorithm: foundations and theory.* Cambridge, MA: MIT press.