
FITNESS UNIFORM SELECTION TO PRESERVE GENETIC DIVERSITY

Marcus Hutter

IDSIA, Galleria 2, CH-6928 Manno-Lugano, Switzerland
marcus@idsia.ch <http://www.idsia.ch/~marcus>

Keywords

Evolutionary algorithms, fitness uniform selection strategy, preserve diversity, local optima, evolution, correlated recombination, crossover.

Abstract

In evolutionary algorithms, the fitness of a population increases with time by mutating and recombining individuals and by a biased selection of more fit individuals. The right selection pressure is critical in ensuring sufficient optimization progress on the one hand and in preserving genetic diversity to be able to escape from local optima on the other. We propose a new selection scheme, which is uniform in the fitness values. It generates selection pressure towards sparsely populated fitness regions, not necessarily towards higher fitness, as is the case for all other selection schemes. We show that the new selection scheme can be much more effective than standard selection schemes.

1 Introduction

Evolutionary algorithms (EA): Evolutionary algorithms are capable of solving complicated optimization tasks in which an objective function $f: I \rightarrow \mathbb{R}$ shall be maximized. $i \in I$ is an individual from the set I of feasible solutions. Infeasible solutions due to constraints may also be considered by reducing f for each violated constraint. A population $P \subseteq I$ of individuals is maintained and updated as follows: one or more individuals are selected according to some selection strategy. In generation based EAs, the selected individuals are recombined (e.g. crossover) and mutated, and constitute the new population. We prefer the more incremental, steady-state population update, which selects (and possibly deletes) only one or two individuals from the current population and adds the newly recombined and mutated individuals to it. We are interested in finding a single individual of maximal objective value f for difficult multimodal and deceptive problems.

Standard selection schemes (STD): The standard selection schemes (abbreviated by STD in the following), proportionate, truncation, ranking and tournament selection all favor individuals of higher fitness [Gol89, GD91, BT95, BT97]. This is also true for less common schemes, like Boltzmann selection [MT93]. The fitness function is identified with the objective function (possibly after a monotone transformation). In linear proportionate selection the probability of selecting an individual depends linearly on its fitness [Hol75]. In truncation selection the $\alpha\%$ fittest individuals are selected, usually with multiplicity $\frac{1}{\alpha\%}$ in order to keep the population size fixed [MSV94]. (Linear) ranking selection orders the individuals according to their fitness. The selection probability is, then, a (linear) function of the rank [Whi89]. Tournament selection [Bak85], which selects the best l out of k individuals has primarily developed for steady-state EAs, but can be adapted to generation based EAs. All these selection schemes have the property (and goal!) to increase the average fitness of a population, i.e. to evolve the population towards higher fitness. For a population with a Gaussian fitness distribution, the probability of selecting an individual and the effect of selection is shown in Figure 1.

The problem of the right selection pressure: The standard selection schemes STD, together with mutation and recombination, evolve the population towards higher fitness. If the selection pressure is too high, the EA gets stuck in a local optimum, since the genetic diversity rapidly decreases. The suboptimal genetic material which might help in finding the global optimum is deleted too rapidly (premature convergence). On the other hand, the selection pressure cannot be chosen arbitrarily low if we want EA to be effective. In difficult optimization problems, suitable population sizes, mutation and recombination rates, and selection parameters, which influence the selection intensity, are usually not known beforehand. Often, constant values are not sufficient at all. There are various suggestions to dynamically determine and adapt the parameters [Esh91, BHS91, Her92, SVM94]. Other approaches to preserve genetic diversity are fitness sharing [GR87], crowding [Jon75] and local mating [CJ91]. They depend on the proper design of a neighborhood function based on the specific problem structure and/or coding.

The main idea: Here, we propose a new selection scheme, based on the insight that we are not primarily interested in a population converging to maximal fitness, but only in a single individual of maximal fitness. The scheme automatically creates a suitable selection pressure and preserves genetic diversity better than STD. The proposed fitness uniform selection scheme FUSS (see also Figure 1) is defined as follows: *if the lowest/highest fitness values in the current population P are $f_{min/max}$, we select a fitness value f uniformly in the interval $[f_{min}, f_{max}]$. Then, the individual $i \in P$ with fitness nearest to f is selected and a copy is added to P , possibly after mutation and recombination.* We will see that FUSS maintains genetic diversity much better than STD, since a distribution over the fitness values is used, unlike STD, which all use a distribution over individuals. Premature convergence is avoided in FUSS by abandoning convergence at all. Nevertheless there is a selection pressure in FUSS towards higher fitness. The probability of selecting a specific individual is proportional to the distance to its nearest fitness neighbor. In a population with a high density of unfit and low density of fit individuals, the fitter ones are effectively favored.

Contents: In *Section 2* we discuss the problems of local optima and exponential takeover [GD91] in STD. Motivated by the need to preserve genetic diversity, we define the fitness uniform selection scheme FUSS. We discuss under which circumstances FUSS leads to an (approximate) fitness uniform population.

Further properties of FUSS are discussed in *Section 3*, especially, how FUSS creates selection pressure towards higher fitness and how it preserves diversity better than STD. Further topics are the equilibrium distribution and the transformation properties of FUSS under linear and non-linear transformations.

In *Section 4* we demonstrate by way of a simple optimization example that an EA with FUSS can optimize much faster than with STD. We show that crossover can be effective in FUSS, even when ineffective in STD. Furthermore, FUSS and STD are compared to random search with and without crossover.

There is a possible slowdown when including recombination, as discussed in *Section 5*, which can be avoided by using a scale independent pair selection. It is a “best” compromise between unrestricted recombination and recombination of individuals with similar fitness only.

To simplify the discussion we have concentrated on the case of discrete, equi-spaced fitness values. In many practical problems, the fitness function is continuously valued. FUSS and some of the discussion of the previous sections is generalized to the continuous case in *Section 6*.

A summary, conclusions and further discussions can be found in *Section 7*.

The focus of this work is on a theoretical analysis of FUSS. Implementation details and numerical results for various test-functions and for real-world problems will be presented elsewhere.

2 Fitness Uniform Selection Strategy (FUSS)

The problem of local optima: Proportionate, truncation, ranking and tournament are the standard (STD) selection algorithms used in evolutionary optimization. They have the following property: if a local optimum i^{lopt} has been found, the number of individuals with fitness $f^{lopt} = f(i^{lopt})$ increases exponentially. Assume a low mutation and recombination rate, or, for instance, truncation selection *after* mutation and recombination. Further, assume that it is very difficult to find an individual more fit than i^{lopt} . The population will then degenerate and will consist mostly of i^{lopt} after a few rounds. This decreased diversity makes it even more unlikely that f^{lopt} gets improved. The suboptimal genetic material which might help in finding the global optimum has been deleted too rapidly. On the other hand, too high mutation and recombination rates convert the EA into an inefficient random search. In the following we suggest a new a new selection scheme, which automatically generates a suitably adapting selection pressure.

The fitness uniform selection scheme (FUSS): For simplicity we start with a fitness function $f : I \rightarrow F$ with discrete equi-spaced values $F = \{f_{min}, f_{min} + \varepsilon, f_{min} + 2\varepsilon, \dots, f_{max} - \varepsilon, f_{max}\}$. The continuous valued case $F = [f_{min}, f_{max}]$ is considered later. The fitness uniform selection scheme (FUSS) is defined as follows: randomly select a fitness value f uniformly from the fitness values F . Randomly (uniformly) select an individual i from population P with fitness f . Add another copy of i to P .

Note the two stage uniform selection process which is very different from a one step uniform selection of an individual of P (see Figure 1). In STD, inertia increases with population size. A large mass of unfit individuals reduces the probability of selecting fit individuals. This is not the case for FUSS. Hence, without loss of performance, we can define a *pure model*, in which no individual is ever deleted; the population size increases with time. No genetic material is ever discarded and no fine-tuning in population size is necessary. What may prevent the pure model from being applied to practical problems are not computation time issues, but memory problems. If space gets a problem we delete individuals from the most occupied fitness levels. Most of the following statements remain valid with this modification.

Asymptotically fitness uniform population: The expected number of individuals per fitness level f after t selections is $n_t(f) = n_0(f) + t/|F|$, where $n_0(f)$ is the initial distribution. Hence, asymptotically the fitness levels get uniformly occupied by a population fraction

$$\frac{n_t(f)}{|P_t(f)|} = \frac{n_0(f) + t/|F|}{|P_0(f)| + t} \rightarrow \frac{1}{|F|} \quad \text{for } t \rightarrow \infty,$$

where $P_t(f)$ is the set of individuals at time t with fitness f .

Fitness gaps and continuous fitness: We made two unrealistic assumptions. First, we assumed that each fitness level is initially occupied. If the smallest/largest fitness

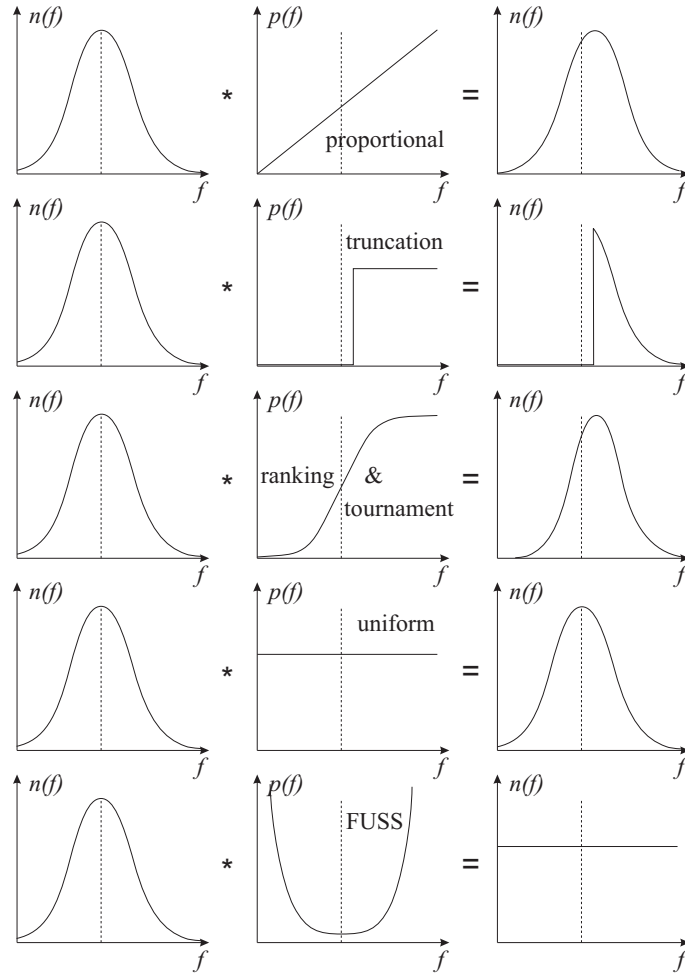


Figure 1: *Effects of proportionate, truncation, ranking & tournament, uniform, and fitness uniform (FUSS) selection on the fitness distribution in a generation based EA. The left/right diagrams depict fitness distributions before/after applying the selection schemes depicted in the middle diagrams.*

values in P_t are $f_{min/max}^t$ we extend the definition of FUSS by selecting a fitness value f uniformly in the interval $[f_{min}^t - \frac{1}{2}\varepsilon, f_{max}^t + \frac{1}{2}\varepsilon]$ and an individual $i \in P_t$ with fitness nearest to f . This also covers the case when there are missing intermediate fitness values, and also works for continuous valued fitness functions ($\varepsilon \rightarrow 0$).

Mutation and recombination: The second assumption was that there is no mutation and recombination. In the presence of small mutation and/or recombination rates eventually each fitness level will become occupied and the occupation fraction is still asymptotically approximately uniform. For larger rates the distribution will be no longer uniform, but the important point is that the occupation fraction of *no* fitness level decreases to zero for $t \rightarrow \infty$, unlike for STD. Furthermore, FUSS selects by construction uniformly in the fitness levels, even if the levels are not uniformly occupied. We will see

that this is the more important property.

3 Properties of Fuss

FUSS effectively favors fit individuals: FUSS preserves diversity better than STD, but the latter have a (higher) selection pressure towards higher fitness, which is necessary for optimization. At first glance it seems that there is no such pressure at all in FUSS, but this is deceiving. As FUSS selects uniformly in the fitness levels, individuals of low populated fitness levels are effectively favored. The probability of selecting a specific individual with fitness f is inverse proportional to $n_t(f)$ (see Figure 1). In a typical (FUSS) population there are many unfit and only a few fit individuals. Hence, fit individuals are effectively favored until the population becomes fitness uniform. Occasionally, a new higher fitness level is discovered and occupied by a new individual, which then, again, is favored.

No takeover in FUSS: With FUSS, takeover of the highest fitness level never happens. The concept of takeover time [GD91] is meaningless for FUSS. The fraction of fittest individuals in a population is always small. This implies that the average population fitness is always much lower than the best fitness. Actually, a large number of fit individuals is usually not the true optimization goal. A single fittest individual usually suffices to having solved the optimization task.

FUSS may also favor unfit individuals: Note, if it is also difficult to find individuals of low fitness, i.e. if there are only few individuals of low fitness, FUSS will also favor these individuals. Half of the time is “wasted” in searching on the wrong end of the fitness scale. This possible slowdown by a factor of 2 is usually acceptable. In Section 4 we will see that in certain circumstances this behaviour can actually speedup the search. In general, fitness levels which are difficult to reach, are favored.

Distribution within a fitness level: Within a fitness level there is no selection pressure which could further exponentially decrease the population in certain regions of the individual space. This (exponential) reduction is the major enemy of diversity, which is suppressed by FUSS. Within a fitness level, the individuals freely drift around (by mutation). Furthermore, there is a steady stream of individuals into and out of a level by (d)evolution from (higher)lower levels. Consequently, FUSS develops an equilibrium distribution which is nowhere zero. We expect FUSS to somewhat lower (but not to solve) the problems associated with genetic drift. The above does also not mean that the distribution within a level is uniform. For instance, if there are two (local) maxima of same height, a very broad one and a very narrow one, the broad one may be populated much more than the narrow one, since it is much easier to “find”.

Steady creation of individuals from every fitness level: In STD, a wrong step (mutation) at some point in evolution might cause further evolution in the wrong direction. Once a local optimum has been found and all unfit individuals were eliminated it

is very difficult to undo the wrong step. In FUSS, all fitness levels remain occupied from which new mutants are steadily created, occasionally one leading to further evolution in a more promising direction.

Transformation properties of FUSS: FUSS (with continuous fitness) is independent of a scaling and a shift of the fitness function, i.e. $\text{FUSS}(\tilde{f})$ with $\tilde{f}(i) := a \cdot f(i) + b$ is identical to $\text{FUSS}(f)$. This is true even for $a < 0$, since FUSS searches for maxima *and* minima, as we have seen. It is not independent of a non-linear (monotone) transformation unlike tournament, ranking and truncation selection. The non-linear transformation properties are more like the ones of proportionate selection.

4 A Simple Example

In the following we compare the performance of fitness uniform selection (FUSS), random search (RAND) and standard selection (STD) with and without recombination on a simple test example. We regard it as a prototype for deceptive multimodal functions. The example should demonstrate why FUSS can be superior to RAND and STD. Numerical results are briefly discussed at the end of the section.

Simple 2D example: Consider individuals $(x, y) \in I := [0, 1] \times [0, 1]$, which are tuples of real numbers, each coordinate in the interval $[0, 1]$. The example models individuals possessing up to 2 “features”. Individual i possesses feature I_1 if $i \in I_1 := [a, a+\Delta] \times [0, 1]$, and feature I_2 if $i \in I_2 := [0, 1] \times [b, b+\Delta]$. The fitness function $f : I \rightarrow \{1, 2, 3\}$ is defined as

$$f(x, y) = \begin{cases} 1 & \text{if } (x, y) \in I_1 \setminus I_2, \\ 2 & \text{if } (x, y) \in I_2 \setminus I_1, \\ 3 & \text{if } (x, y) \notin I_1 \cup I_2, \\ 4 & \text{if } (x, y) \in I_1 \cap I_2. \end{cases}$$

We assume $\Delta \ll 1$. Individuals with none of the two features ($i \in I \setminus (I_1 \cup I_2)$) have fitness $f = 3$. These “local $f = 3$ optima” occupy most of the individual space I , namely a fraction $(1 - \Delta)^2$. It is disadvantageous for an individual to possess only one of the two features ($i \in (I_1 \setminus I_2) \cup (I_2 \setminus I_1)$), since $f = 1/f = 2$ in this case. In combination ($i \in I_1 \cap I_2$), the two features lead to the highest fitness, but the global maximum $f = 4$ occupies the smallest fraction Δ^2 of the individual space I . With a fraction $\Delta(1 - \Delta)$, the $f = 1/f = 2$ minima are in between.

Random search: Individuals are created uniformly in the unit square. The “local optimum” $f = 3$ is easy to “find”, since it occupies nearly the whole space. The global optimum $f = 4$ is difficult to find, since it occupies only $\Delta^2 \ll 1$ of the space. The expected time, i.e. the expected number of individuals created and tested until one with $f = 4$ is found, is $T_{\text{RAND}} = \frac{1}{\Delta^2}$. Here and in the following, the “time” T is defined as the

number of created individuals until the *first* optimal individual (with $f = 4$) is found. T is neither a takeover time nor the number of generations (we consider steady-state EAs).

Random search with crossover: Let us occasionally perform a recombination of individuals in the current population. We combine the x -coordinate of one uniformly selected individual i_1 with the y coordinate of another individual i_2 . This crossover operation maintains a uniform distribution of individuals in $[0, 1]^2$. It leads to the global optimum if $i_1 \in I_1$ and $i_2 \in I_2$. The probability of selecting an individual in I_i is $\Delta(1 - \Delta) \approx \Delta$ (we assumed that the global optimum has not yet been found). Hence, the probability that I_1 crosses with I_2 is Δ^2 . The time to find the global optimum by random search including crossover is still $\sim \frac{1}{\Delta^2}$.

Mutation: The result remains valid (to leading order in $\frac{1}{\Delta}$) if, instead of a random search, we uniformly select an individual and mutate it according to some probabilistic, sufficiently mixing rule, which preserves uniformity in $[0, 1]$. One popular such mutation operator is to use a sufficiently long binary representation of each coordinate, like in genetic algorithms, and flip a single bit. In the following, we assume a mutation operator which mutates with probability $\frac{1}{2}/\frac{1}{2}$ the first/second coordinate, preserves uniformity, is sufficiently mixing, and leaves the other coordinate unchanged, like the single-bit-flip operator.

Standard selection with crossover: The $f = 1$ and $f = 2$ individuals contain useful building blocks, which could speedup the search by a suitable selection and crossover scheme. Unfortunately, the standard selection schemes favor individuals of higher fitness and will diminish the $f = 1/f = 2$ population fraction. The probability of selecting $f = 1/f = 2$ individuals is even smaller than in random search. Hence $T_{STD} \sim \frac{1}{\Delta^2}$. Standard selection does not improve performance, even not in combination with crossover, although crossover is well suited to produce the needed recombination.

FUSS: At the beginning, only the $f = 3$ level is occupied and individuals are uniformly selected and mutated. The expected time till an $f = 1/f = 2$ individual in $I_1 \cup I_2$ is created is $T_1 \approx \frac{1}{\Delta}$ (not $\frac{1}{2\Delta}$, since only one coordinate is mutated). From this time on FUSS will select one half(!) of the time the $f = 1/f = 2$ individual(s) and only the remaining half the abundant $f = 3$ individuals. When level $f = 1$ and level $f = 2$ are occupied, the selection probability is $\frac{1}{3} + \frac{1}{3}$ for these levels. With probability $\frac{1}{2}$ the mutation operator will mutate the y coordinate of an individual in I_1 or the x coordinate of an individual in I_2 and produces a new $f = 1/2/4$ individual. The relative probability of creating an $f = 4$ individual is Δ . The expected time to find this global optimum from the $f = 1/f = 2$ individuals, hence, is $T_2 = [(\frac{1}{2} \dots \frac{2}{3}) \times \frac{1}{2} \times \Delta]^{-1}$. The total expected time is $T_{FUSS} \approx T_1 + T_2 = \frac{4}{\Delta} \dots \frac{5}{\Delta} \ll \frac{1}{\Delta^2} \sim T_{STD}$. FUSS is much faster by exploiting unfit $f = 1/f = 2$ individuals. This is an example where (local) minima can help the search. Examples where a low local maxima can help in finding the global maximum, but where standard selection sweeps over too quickly to higher but useless local maxima, can also be constructed.

FUSS with crossover: The expected time till an $f = 1$ individual in I_1 and an $f = 2$ individual in I_2 is found is $T_1 \approx \frac{1}{\Delta}$, even with crossover. The probability of selecting an $f = 1/f = 2$ individual is $\frac{1}{3}/\frac{1}{3}$. Thus, the probability that a crossing operation crosses I_1 with I_2 is $(\frac{1}{3})^2$. The expected time to find the global optimum from the $f = 1/f = 2$ individuals, hence, is $T_2 = 9 \cdot O(1)$, where the $O(1)$ factor depends on the frequency of crossover operations. This is far faster than by STD, even if the $f = 1/f = 2$ levels were local maxima, since to get a high standard selection probability, the level has first to be taken over, which itself needs some time depending on the population size. In FUSS a single $f = 1$ and a single $f = 2$ individual suffice to guarantee a high selection probability and an effective crossover. Crossover does not significantly decrease the *total* time $T_{FUSSX} \approx T_1 + T_2 = \frac{1}{\Delta} + O(9)$, but for a suitable 3D generalization we get a large speedup by a factor of $\frac{1}{\Delta}$.

Simple 3D example: We generalize the 2D example to D-dimensional individuals $\mathbf{x} \in [0, 1]^D$ and a fitness function

$$f(\mathbf{x}) := (D+1) \cdot \prod_{d=1}^D \chi_d(\mathbf{x}) - \max_{1 \leq d \leq D} d \cdot \chi_d(\mathbf{x}) + d + 1,$$

where $\chi_d(\mathbf{x})$ is the characteristic function of feature I_d

$$\chi_d(\mathbf{x}) := \begin{cases} 1 & \text{if } a_i \leq x_i \leq a_i + \Delta, \\ 0 & \text{else.} \end{cases}$$

For $D = 2$, f coincides with the 2D example. For $D = 3$, the fractions of $[0, 1]^3$ where $f = 1/2/3/4/5$ are approximately $\Delta^2/\Delta^2/\Delta^2/1/\Delta^3$. With the same line of reasoning we get the following expected search times for the global optimum:

$$T_{RAND} \sim T_{STD} \sim \frac{1}{\Delta^3},$$

$$T_{FUSS} \sim \frac{1}{\Delta^2}, \quad T_{FUSSX} \sim \frac{1}{\Delta}.$$

This demonstrates the existence of problems, where FUSS is much faster than RAND and STD, and that crossover can give a further boost in FUSS, even when ineffective in combination with STD.

Numerical results: An EA with FUSS and STD has been implemented. First experiments confirm the superiority of FUSS also for other complicated multimodal and deceptive fitness functions. The asymptotic behavior of the convergence times T for $\Delta \rightarrow 0$ for the previous example has been verified. We got similar results for the function $f(x+a, y+b) = 2\frac{x^2}{x^2+\Delta^2} + \frac{y^2}{y^2+\Delta^2} + 4e^{-(x^2+y^2)/\Delta^2}$, which is a continuous version of the 2D example. We further applied FUSS to the Traveling Salesman Problem. We considered $10^{1..3}$ cities with random matrix distances, random initial paths, random 1-Opt and 2.5-Opt mutation operators, inverse length as fitness, but no crossover yet. The

solutions found by FUSS are consistently and significantly better than those found by STD. The current implementation can in no way compete with up-to-date TSP-solvers [JM97, ACR00], but this was not the intention of the comparison. The results are very preliminary yet. Implementation details and detailed simulation results will be presented in a forthcoming publication.

5 Recombination

Worst case analysis without recombination: We now want to estimate the maximal possible slowdown of FUSS compared to STD. Let us assume that all individuals in STD have fitness f , and once one individual with fitness $f+\varepsilon$ has been found, takeover of level $f+\varepsilon$ is quick. Let us assume that this quick takeover is actually good (e.g. if there are no local maxima). The selection probability of individuals of same fitness is equal. For FUSS we assume individuals in the range of f_{min} and f . Uniformity is *not* necessary. In the worst case, a selection of an individual of fitness $< f$ never leads to an individual of fitness $\geq f$, i.e. is always useless. The probability of selecting an individual with fitness f is $\geq \frac{1}{|F|}$. At least every $|F|$ th FUSS selection corresponds to a STD selection. Hence, we expect a maximal slowdown by a factor of $|F|$, since FUSS “simulates” STD statistically every $|F|$ th selection. It is possible to artificially construct problems where this slowdown occurs (unimodal function, local mutation $x \rightarrow x \pm \varepsilon$, no crossover). We have never observed this large slowdown in our experiments. For the more complicated multimodal and deceptive objective functions we were interested in, FUSS often outperformed STD in solution quality *and* time.

Quadratic slowdown due to recombination: We have seen that $T_{FUSS} \leq |F| \cdot T_{STD}$. In the presence of recombination, a *pair* of individuals has to be selected. The probability that FUSS selects *two* individuals with fitness f is $\geq \frac{1}{|F|^2}$. Hence, in the worst case, there could be a slowdown by a factor of $|F|^2$ — for *independent* selection we expect $T_{FUSS} \leq |F|^2 \cdot T_{STD}$. This potential quadratic slowdown can be avoided by selecting one fitness value at random, and then two individuals of this single fitness value. For this *dependent* selection, we expect $T_{FUSS} \leq |F| \cdot T_{STD}$. One the other hand, crossing two individuals of different fitness can also be advantageous, like the crossing of $f=1$ with $f=2$ individuals in the 2D example of Section 4.

Scale independent pair selection: It is possible to (nearly) have the best of independent and dependent selection: a high selection probability $p(f, f') \sim \frac{1}{|F|}$ if $f \approx f'$ and $p(f, f') \sim \frac{1}{|F|^2}$ otherwise, with uniform marginal $p(f) = \frac{1}{|F|}$. The idea is to use a strongly correlated joint distribution for selecting a fitness pair. A “scale independent” probability distribution $p(f, f') \sim \frac{1}{|f-f'|}$ is appropriate. We define the joint probability $\tilde{p}(f, f')$ of selecting two individuals of fitness f and f' and the marginal $\tilde{p}(f)$ as

$$\tilde{p}(f, f') := \frac{1}{2|F| \ln |F|} \cdot \frac{1}{\frac{1}{\varepsilon}|f-f'| + 1}, \quad (1)$$

$$\tilde{p}(f) := \sum_{f' \in F} \tilde{p}(f, f') = \sum_{f' \in F} \tilde{p}(f', f).$$

We assume $|F| \geq 3$ in the following. The $+1$ in the denominator has been added to regularize the expression for $f = f'$. The factor $(2|F| \ln |F|)^{-1}$ ensures correct normalization for $|F| \rightarrow \infty$. More precisely, using $\ln \frac{b+1}{a} \leq \sum_{i=a}^b \frac{1}{i} \leq \ln \frac{b}{a-1}$, one can show that

$$1 - \frac{1}{\ln |F|} \leq \sum_{f, f' \in F} \tilde{p}(f, f') \leq 1, \quad \frac{1}{2} \leq |F| \cdot \tilde{p}(f) \leq 1$$

i.e. \tilde{p} is not strictly normalized to 1 and the marginal $\tilde{p}(f)$ is only approximately (within a factor of 2) uniform. The first defect can be corrected by appropriately increasing the diagonal probabilities $\tilde{p}(f, f)$. This also solves the second problem.

$$p(f, f') := \begin{cases} \tilde{p}(f, f') & \text{for } f \neq f' \\ \tilde{p}(f, f') + [\frac{1}{|F|} - \tilde{p}(f)] & \text{for } f = f' \end{cases} \quad (2)$$

Properties of $p(f, f')$: p is normalized to 1 with uniform marginal

$$p(f) := \sum_{f' \in F} p(f, f') = \frac{1}{|F|}, \quad p(f, f') \geq \tilde{p}(f, f'),$$

$$\sum_{f, f' \in F} p(f, f') = \sum_{f \in F} p(f) = 1,$$

Apart from a minor additional logarithmic suppression of order $\ln |F|$ we have the desired behaviour $p(f, f') \sim \frac{1}{|F|}$ for $f \approx f'$ and $p(f, f') \sim \frac{1}{|F|^2}$ otherwise:

$$p(f, f \pm m\varepsilon) \geq \frac{1}{2 \ln |F|} \cdot \frac{1}{m+1} \cdot \frac{1}{|F|},$$

$$p(f, f') \geq \frac{1}{2 \ln |F|} \cdot \frac{1}{|F|^2}.$$

During optimization, the minimal/maximal fitness of an individual in population P_t is $f_{min/max}^t$. In the definition of p one has to use $F_t := \{f_{min}^t, f_{min}^t + \varepsilon, \dots, f_{max}^t\}$ instead of F , i.e. $|F| \leftrightarrow |F_t| = \frac{1}{\varepsilon}(f_{max}^t - f_{min}^t) + 1 \leq |F|$.

6 Continuous Fitness Functions

Effective discretization scale: Up to now we have considered a discrete valued fitness function with values in $F = \{f_{min}, f_{min} + \varepsilon, \dots, f_{max}\}$. In many practical problems, the fitness function is continuous valued with $F = [f_{min}, f_{max}]$. We generalize FUSS, and some of the discussion of the previous sections to the continuous case by replacing the discretization scale ε by an effective (time-dependent) discretization scale $\hat{\varepsilon}$. By

construction, FUSS shifts the population towards a more uniform one. Although the fitness values are no longer equi-spaced, they still form a discrete set for finite population P . For a fitness uniform distribution, the average distance between (fitness) neighboring individuals is $\frac{1}{|P_t|-1}(f_{max}^t - f_{min}^t) =: \hat{\varepsilon}$. We define $\hat{F}_t := \{f_{min}^t, f_{min}^t + \hat{\varepsilon}, \dots, f_{max}^t\}$. $|\hat{F}_t| = \frac{1}{\hat{\varepsilon}}(f_{max}^t - f_{min}^t) + 1 = |P_t|$.

FUSS: Fitness uniform selection for a continuous valued function has already been mentioned in Section 2. We just take a uniform random fitness f in the interval $[f_{min}^t - \frac{1}{2}\hat{\varepsilon}, f_{max}^t + \frac{1}{2}\hat{\varepsilon}]$. One may even take the limit $\hat{\varepsilon} \rightarrow 0$ in this case probably without harming FUSS. Independent and dependent fitness pair selection as described in the last section works analogously. An $\hat{\varepsilon}=0$ version of correlated selection does not exist; a non-zero $\hat{\varepsilon}$ is important. A discrete pair (f, f') is drawn with probability $p(f, f')$ as defined in (1) and (2) with ε and F replaced by $\hat{\varepsilon}$ and \hat{F}_t . The additional suppression $\ln |\hat{F}_t| = \ln |P_t|$ is small for all practical population sizes.

In all cases an individual with fitness nearest to f (f') is selected from the population P (randomly if there is more than one nearest individual).

Discussion: If we assume a fitness uniform distribution, a worst case bound $T_{FUSS} \lesssim \sum_{t=1}^{T_{STD}} |P_t|$ seems plausible, since the probability of selecting the best individual is approximately $|P_t|$. For constant population size we get a bound $T_{FUSS} \lesssim |P| \cdot T_{STD}$. For the preferred non-deletion case with population size $|P_t| = t$ the bound gets much worse $T_{FUSS} \lesssim \frac{1}{2} T_{STD}^2$. This possible (but not necessary!) slowdown has similarities to the slowdown problems of proportionate selection in later optimization stages. Larger choices of $\hat{\varepsilon}$ may be favorable if the standard choice causes problems.

7 Summary & Conclusions

We have addressed the problem of balancing the selection intensity in EAs, which determines speed versus quality of a solution. We invented a new fitness uniform selection scheme FUSS. It generates a selection pressure towards sparsely populated fitness levels. This property is unique to FUSS as compared to other selection schemes (STD). It results in the desired high selection pressure towards higher fitness if there are only a few fit individuals. The selection pressure is automatically reduced when the number of fit individuals increases. A joint pair selection scheme for recombination has been defined, but not yet implemented. A heuristic worst case analysis of FUSS compared to STD has been given. FUSS solves the problem of exponential takeover and the resulting loss of genetic diversity of STD, while still generating enough selection pressure. It does not help in getting a more uniform distribution within a fitness level. We showed analytically by way of a simple example that FUSS can be much more effective than STD. FUSS should be compared to STD on other problems to further explore its efficacy and limitations. First results look encouraging. Of special interest is whether FUSS could improve up-to-date EAs that solve difficult combinatoric optimization problems,

like TSPs. We expect FUSS to be superior to STD in cases where an EA with STD effectively gets trapped into local optima.

References

- [ACR00] D. Applegate, W. Cook, and A. Rohe. Chained Lin-Kernighan for large traveling salesman problems. Technical report, 2000. http://www.keck.caam.rice.edu/reports/chained_lk.ps.
- [Bak85] J. E. Baker. Adaptive selection methods for genetic algorithms. In J. J. Grefenstette, editor, *Proceedings of the 1st International Conference on Genetic Algorithms and their Applications*, pages 101–111, Pittsburgh, PA, 1985. Lawrence Erlbaum Associates.
- [BHS91] T. Bäck, F. Hoffmeister, and H. P. Schwefel. A survey of evolution strategies. In R. K. Belew and L. B. Booker, editors, *Proceedings of the 4th International Conference on Genetic Algorithms*, pages 2–9, San Diego, CA, July 1991. Morgan Kaufmann.
- [BT95] T. Blicke and L. Thiele. A mathematical analysis of tournament selection. In L. J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms (ICGA'95)*, pages 9–16, San Francisco, California, 1995. Morgan Kaufmann Publishers.
- [BT97] T. Blicke and L. Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4(4):361–394, 1997.
- [CJ91] R. J. Collins and D. R. Jefferson. Selection in massively parallel genetic algorithms. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, San Mateo, CA, 1991. Morgan Kaufmann Publishers.
- [Esh91] L. J. Eshelman. The CHC adaptive search algorithm: How to safe search when engaging in nontraditional genetic recombination. In G. J. E. Rawlings, editor, *Foundations of genetic algorithms*, pages 265–283. Morgan Kaufmann, San Mateo, 1991.
- [GD91] D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. In G. J. E. Rawlings, editor, *Foundations of genetic algorithms*, pages 69–93. Morgan Kaufmann, San Mateo, 1991.
- [Gol89] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, Mass., 1989.

- [GR87] D. E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodalfunction optimization. In J. J. Grefenstette, editor, *Proceedings of the 2nd International Conference on Genetic Algorithms and their Applications*, pages 41–49, Cambridge, MA, July 1987. Lawrence Erlbaum Associates.
- [Her92] M. Herdy. Reproductive isolation as strategy parameter in hierarchically organized evolution strategies. In R. Männer and B. Manderick, editors, *Parallel problem solving from nature 2*, pages 207–217, Amsterdam, 1992. North-Holland.
- [Hol75] John H. Holland. *Adpatation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, 1975.
- [JM97] D. S. Johnson and A. McGeoch. The traveling salesman problem: A case study. In E. H. L. Aarts and J. K. Lenstra, editors, *Local Search in Combinatorial Optimization*, Discrete Mathematics and Optimization, chapter 8, pages 215–310. Wiley-Interscience, Chichester, England, 1997.
- [Jon75] K. de Jong. An analysis of the behavior of a class of genetic adaptive systems. *Dissertation Abstracts International*, 36(10), 5140B, 1975.
- [MSV94] Heinz Mühlenbein and Dirk Schlierkamp-Voosen. The science of breeding and its application to the breeder genetic algorithm (BGA). *Evolutionary Computation*, 1(4):335–360, 1994.
- [MT93] M. de la Maza and B. Tidor. An analysis of selection procedures with particular attention paid to proportional and Boltzmann selection. In S. Forrest, editor, *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 124–131, San Mateo, CA, USA, 1993. Morgan Kaufmann.
- [SVM94] D. Schlierkamp-Voosen and H. Mühlenbein. Strategy adaptation by competing subpopulations. In Y. Davidor, H.-P. Schwefel, and R. Männer, editors, *Parallel Problem Solving from Nature – PPSN III*, pages 199–208, Berlin, 1994. Springer. Lecture Notes in Computer Science 866.
- [Whi89] D. Whitley. The GENITOR algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms (ICGA '89)*, pages 116–123, San Mateo, California, 1989. Morgan Kaufmann Publishers, Inc.