

- [33] Sumida, B.H. & Hamilton, W.D. (1994), "Both Wrightian and 'Parasite' Peak Shifts Enhance Genetic Algorithm Performance in the Travelling Salesman Problem", in Paton, R.C. (ed) *Computing with Biological Metaphors*, Chapman and Hall.
- [34] Voigt, H-M., Santibanez-Koref, I., Born, J. (1992), "Hierarchically Structured Distributed GAs", in Manner, R. & Manderick, B. (eds), *Parallel Problem Solving from Nature*, 2, Elsevier: Amsterdam, 155-164.

## Exploiting mate choice in evolutionary computation: Sexual selection as a process of search, optimization, and diversification

Geoffrey F. Miller  
Cognitive and Computing Sciences  
University of Sussex  
Falmer, Brighton BN1 9QH, England  
geoffm@cogs.susx.ac.uk

### Abstract

Sexual selection through mate choice is a powerful evolutionary process that has been important in the success of sexually-reproducing animals and flowering plants. Over the short term, mate preferences evolve because they improve the outcome of sexual recombination. Over the long term, assortative mate preferences can help maintain genetic diversity, promote speciation, and facilitate evolutionary search through optimal outbreeding; selective mate preferences can reinforce the speed, accuracy, and efficiency of natural selection, can foster the discovery and propagation of evolutionary innovations, and can function as aesthetic selection criteria. These strengths of sexual selection complement those of natural selection, so using both together may prove particularly fruitful in evolutionary computation. This paper reviews the biological theory of sexual selection and some possible applications of sexual selection in evolutionary search, optimization, and diversification. Simulation results are used to illustrate some key points.

### 1 Introduction: The Evolutionary Importance of Sexual Selection

The overwhelming biological success of sexually-reproducing animals and flowering plants has often been attributed to the raw power of sexual recombination. Yet the diagnostic feature shared by these two groups is not just sexual recombination *per se* (which bacteria and non-flowering plants also use), but rather, sexual selection through mate choice. Typically, animals are sexually-selected by opposite-sex conspecifics (Darwin, 1871) and flowering plants are sexually selected by pollinators such as insects and hummingbirds (Darwin, 1862). This suggests that the evolution of phenotypic complexity and diversity may be driven not simply by natural-selective adaptation to niches, but by a complementary interplay of natural selection and sexual selection. If mate choice has been instrumental in the evolutionary success of higher animals and plants, perhaps it has been under-estimated as a process of search, optimization, and diversification.

Mate preferences evolve because they improve the outcome of sexual recombination for the organisms that use them. Selective or assortative mating based on mate preferences tends to produce offspring with higher viability, fertility, or attractiveness than random mating. Whereas the merits of recombination through random mating are so unclear that biologists still argue vehemently about why sex evolved

(see Michod & Levin, 1988; Ridley, 1993), the merits of selective mate choice are now almost universally recognized (see Andersson, 1994; Cronin, 1992). Indeed, whenever a species uses sexual recombination, and has the sensory-motor capacity for mate choice, mate choice almost always evolves to guide who recombines with whom.

Biological interest in sexual selection has grown enormously in the last 15 years, but has not yet been integrated with the evolutionary computation view of evolution as a process of search and optimization. For example, genetic algorithms typically use recombination (crossover) without mate choice. Crossover can combine useful schemata or building blocks from different parents, and this is the major advantage that genetic algorithms have over other stochastic, population-based, hill-climbing search algorithms (Eshelman & Schaffer, 1993). But recombination has two basic drawbacks under random mating: good schemata can be disrupted, and genetic diversity can be eroded. Mate choice, particularly assortative mating, is the major way that sexually-reproducing organisms lower these costs while preserving recombination's benefits.

This paper aims to promote interest in sexual selection theory, to suggest some practical ways that mate choice may improve evolutionary computation, and to inspire further work in this area. Space limitations preclude a complete discussion of relevant theory, methods, and results. For a longer review of sexual selection theory and its computational implications, see Miller and Todd (in press). For details of simulation methods and results see Todd and Miller (1991) on assortative mating and speciation, Miller and Todd (1993) on selective mating and stochastic runaway effects, and Todd and Miller (1993) on the evolution of parental sexual imprinting under sexual selection. For a discussion of the role of sexual selection in human evolution see Miller (1993; in press) and Ridley (1993); for a comprehensive, up-to-date biological review of sexual selection, see Andersson (1994).

## 2 Why Mate Preferences Evolve

Organisms that reproduce sexually should avoid mating randomly, because the genetic quality of one's mate will determine half the genetic quality of one's offspring (see Pomiankowski, 1988). Assortative mating with an organism similar to oneself can improve offspring fitness by providing them with co-adapted, strongly-linked genes that function well together, and selective mating with high-fitness organisms can improve offspring fitness by giving them genes that have clearly prospered in the current environment. The key to choosing mates adaptively is to evolve a mate choice mechanism that has 'internalized' the likely long-term fitness consequences of reproducing with different kinds of potential mates. The attractiveness of a potential mate should reflect the expected fitness of any offspring that one might have with it.

A major strength of evolutionary computation methods is that they can find solutions without using problem-specific knowledge, algorithms, or heuristics. However, if problem-specific methods could be evolved automatically, they would almost always be useful. Mate preferences can be viewed as evolved heuristics that improve

the ways selection and recombination operate. The evolution of mate preferences then can be seen as a way of automatically incorporating knowledge about the current and past fitness landscape into evolutionary search. Just as evolutionary reinforcement learning (ERL) can 'internalize' utility functions from the environment into an organism's reinforcement learning system (Litman and Ackley, 1991), sexual selection can internalize fitness functions from the environment into an organism's mate choice system. At the level of an evolving species, mate preferences make a species more 'intelligent' by guiding future evolution based on past information about adaptive success and failure.

Evolvable mate preferences lead to complex evolutionary dynamics, because sexually-selected traits adapt to the current distribution of preferences in the population, but preferences also adapt to the current distribution and fitnesses of traits in the population. Over the short term, traits adapt to preferences and vice-versa; but over the long term, both traits and preferences can co-evolve in unpredictable directions (Miller & Todd, 1993). The following section describes some simple ways of incorporating mate choice into evolutionary computation methods.

## 3 Adding Mate Preferences to Genetic Algorithms

Traditional genetic algorithms use a fitness function to select parents and then pairs parents at random for sexual recombination. Sexual selection based on mate preferences requires an additional selection step after parents have been picked. The sexual selection process must somehow allow each parent to sample a number of potential mates and to 'choose' the most acceptable one based on its mate preferences, perhaps stochastically, and perhaps mutually. Although this section focuses on mate preferences for genetic algorithms, similar methods could be used in genetic programming, evolution strategies, classifier systems, and artificial life systems.

Mate preferences can be represented as probability-of-mating (POM) functions defined across an entire  $n$ -dimensional phenotype space (Todd & Miller, 1991; Miller & Todd, 1993). The POM function assigns to every possible phenotypic location, and thus to every possible mate, a certain per-encounter probability of mating, which can be represented as a height in the  $n+1$ st dimension. Peaks on each individual's POM function correspond to ideal, totally attractive mates; low points correspond to repulsive, totally unacceptable mates. Sexual selection is driven by the topography of evolving POM functions interacting with the evolving frequency distribution of available mates in phenotype space, and with the natural-selective fitness landscape.

Mate preferences can be genetically encoded by preference genes that specify the key parameters of POM functions. One way to focus evolution on an efficiently low number of POM parameters is to construct fairly standardized POMs with respect to a fairly limited class of "reference position" in phenotype space. Reference positions could be genetically specified such that they are set to one's own phenotype

(as in assortative mating), one's parent's phenotype (as in parental imprinting), the current population average phenotype (as in human preferences for face shape), or an absolute position in genotype or phenotype space. Given the reference position, the remainder of the POM can be constructed around it, either as a radially symmetric, 'non-directional' function centered on the reference position (as in Todd & Miller, 1991), or as a vector-like 'directional' function that points away from the reference position in some direction, e.g. such that one prefers a mate much larger or more colorful than oneself (as in Miller & Todd, 1993). We have found that sexual selection works better if the total volume under POM functions is normalized to be equal across individuals; otherwise, many mating schemes will favor 'promiscuous' individuals who are attracted to everybody, thereby undermining the effects of sexual selection.

In nature, mate preferences must rely on observable phenotypic information. But simulation could allow individuals to choose mates based on a much richer database: complete and accurate information about their genotype, their phenotype, their past behavior, their performance on fitness tests, their ancestry, or even the expected performance of offspring produced with them. A preference based on direct genotypic information, for example, could be coded in a single gene specifying an ideal Hamming distance between one's own genotype and the genotype of a potential mate. The next section discusses the benefits of this type of assortative mating.

#### 4 Assortative Mating I: Diversification Through Automatic Niching and Speciation

Genetic algorithms are good at hill-climbing and at overcoming some kinds of 'deception', but they have trouble escaping from local optima and finding multiple global optima in complex fitness landscapes. Short-term selection tends to reduce genetic variation through 'exploitation' of local optima, but the success of long-term evolution through 'exploration' depends on maintaining significant genetic variation. Mutation is not enough, because it produces only superficial variants tightly clustered around currently common genotypes, and sexual recombination with random mating is even worse, because given diverse parents it tends to produce lower-variance offspring with intermediate values on quantitative genetic traits. The problem is to preserve 'deep' genetic diversity as opposed to the superficial variation maintained by mutation and recombination. Two solutions have been developed for maintaining deep diversity: niching methods and spatially structured populations. They will be reviewed briefly, and then a new, complementary method for preserving deep diversity will be discussed: assortative mating.

Niching methods maintain deep diversity by distributing sub-populations across multiple fitness peaks. Normally this is difficult because genetic drift tends to disrupt the numerical balance between sub-populations, often leading populations to 'collapse' down to a single peak. Some methods fight genetic drift fairly passively, through replacing individuals with similar offspring somehow: in preselection (Cavichio, 1970; see Mahfoud, 1992), offspring can only replace one of their parents; in the crowding method (De Jong, 1975), offspring replace the genotype most similar in

Hamming distance, selected from a randomly drawn subpopulation of a certain size. Other methods fight drift more actively through giving higher fitness to individuals in uncrowded regions of genotype or phenotype space. In the sharing scheme (Deb & Goldberg, 1989; Goldberg & Richardson, 1987), individuals within a certain genetic or phenotypic distance of each other must divide up the 'locally available' fitness. Sharing is like ecological competition: it tends to distribute individuals across fitness peaks in proportion to the heights of the peaks. Of these niching methods, the active ones work better than the passive ones, and of the active ones, phenotypic sharing seems to work better than genotypic sharing (Goldberg, Deb, & Horn, 1992; Goldberg & Richardson, 1987). However, these methods require the programmer to set various parameters rather carefully, based on assumptions about the number and distribution of fitness peaks.

Spatially structured populations maintain genetic diversity by maintaining geographic diversity, typically through local competition, local mating, local replacement, and fairly low migration rates. Such populations are naturally implemented on parallel computers, so the software method of spatial structuring has often been conflated with the hardware implementation of 'parallel GAs', through they are conceptually distinct (Gordon & Whitney, 1993). In 'coarse-grained' methods inspired by Wright's shifting balance theory (1932), the population is divided up into 'demes' with some degree of migration between them in each generation; demes can be spatially structured in relation to one another according to some topology and distance metric (as in 'stepping stone' models) or can simply exist as nodes equidistant from one another (as in 'island models'). In 'fine-grained' methods (a.k.a. 'diffusion', 'isolation-by-distance', and 'cellular' methods) individuals rather than demes are spatially arranged, typically as nodes on a 2-D grid. Spatial structuring has proven very useful in avoiding premature convergence, maintaining genetic diversity, allowing niche differentiation, finding global optima faster, finding multiple optima in the same run, and exploiting the power of parallel computers (e.g. Collins & Jefferson, 1991; Davidor et al., 1993; Gorges-Schleuter, 1989, 1992; Manderick & Spiessens, 1989; Mühlbein, 1989; 1992). In spatial structuring, niching emerges somewhat passively, because different geographic areas may evolve towards different peaks due simply to different stochastic effects; nevertheless, spatial structuring reliably and efficiently explores multimodal fitness landscapes, and requires fewer assumptions than non-spatial niching methods.

For all the success of niching and spatial structuring in preserving deep genetic diversity, they overlook the main methods that nature uses in generating biodiversity: assortative mating and speciation. In assortative mating, animals pair up based on their similarity, such that "like mates with like". In speciation, a lineage splits apart into reproductively isolated populations (species) that can no longer interbreed. Speciation can be viewed as the extreme outcome of assortative mating. Speciation is important in evolutionary computation because it lets lineages specialize in exploring different peaks in phenotype space, without wasting effort on trying to interbreed across fitness valleys. In nature, speciation creates biodiversity to fill niches; in evolutionary computation, speciation creates separate lineages doing parallel searches in fitness landscapes.

Speciation can be initiated by geographic isolation (Mayr, 1942) or ecological specialization (Dobzhansky, 1937), both of which tend to split populations apart. But for the initial split to be consolidated through reproductive isolation, mate preferences must reinforce the differentiation. Indeed, in Todd and Miller (1991) we showed that mate preferences can spontaneously differentiate in a population (through a combination of genetic drift, mutation pressure, and assortative mating), such that speciation can occur even without geographic isolation or ecological specialization. Our model of 'spontaneous sympatric speciation' through assortative mating is consistent with the fact that sexually-reproducing animals and flowering plants show vastly greater biodiversity and higher speciation rates than other kinds of organisms that do not use mate choice (Miller & Todd, in press). Of course, once speciation occurs, the well-known ecological 'exclusion principle' (two species cannot occupy the same niche in the same area at the same time) tends to push species into separate ecotones or separate habitats. Thus, reproductive isolation through the spontaneous (stochastic) differentiation of mate preferences can initiate and consolidate other adaptive differences between populations, and can lead to ever-increasing biodiversity and niche differentiation.

Some genetic algorithms research has flirted with speciation. Booker (1985) explored the use of evolvable 'mating templates' that allow assortative and selective mating. Some sharing schemes use assortative 'mating restrictions' to minimize the disruptive effects of recombination across different fitness peaks; and it has been found that sharing plus assortative mating works better than sharing alone (Deb & Goldberg, 1989; Goldberg & Richardson, 1987). Still, a more systematic approach to maintaining genetic diversity through assortative mating might prove valuable. For example, mate preferences would be expected to register when a population has split apart on different fitness peaks, by evolving to discourage cross-peak mating (since preferences that encourage cross-peak mating get passed on to inviable hybrids and die out.) Evolvable preferences thus should be able to adaptively and spontaneously promote niching as a lineage evolves, without the programmer having to estimate the number and distribution of peaks ahead of time.

The main problem with maintaining speciation purely through assortative mating is that relative species numbers are subject to genetic drift if competition and replacement are handled as global operations: incipient species tend to disappear due to sampling error. We have developed a method for minimizing the effects of drift under assortative mating through allowing only one parent at a time to exert mate choice (Todd & Miller, 1991), but this only delays the inevitable extinction of species due to drift. Inman Harvey and I have recently been exploring the effects of assortative mating for adaptation in Kauffman's (1993) N-K model of fitness landscapes; so far, this problem of genetic drift seems to override any benefit of assortative mating, at least for a rough sample of K values between 2 and 5, N values of 10 to 100, population sizes of 50 to 500, and mutation rates from 0.3 to 3.0 per genome.

In nature, relative species numbers can be maintained because species diverge ecologically or geographically. In evolutionary computation then, it probably

makes sense to combine ecological niching, spatial structuring, and assortative mating. GA methods could be developed to consolidate speciation by recognizing ecologically differentiated populations through phenotypic cluster analysis, and/or reproductively isolated populations through analyzing patterns of assortative mating, and then allowing each population to run autonomously as a separate GA 'lineage'. Such a method could be implemented on a parallel computer by starting with one species running on one processor, and assigning new species to other processors as they split off, up to the biodiversity limit imposed by the number of processors. Evolved mechanisms for assortative mating should prove critical to the success of such methods because it is the clearest way to allow a population to decide by itself when it 'wants' to speciate.

## 5 Assortative Mating II: Search Through Optimal Outbreeding

Aside from maintaining genetic diversity, another fundamental problem for GAs is maximizing the recombination of useful schemata and building blocks while minimizing the disruption of useful ones. In nature, this seems to be accomplished by 'optimal outbreeding' (see Bateson, 1983), a very common form of assortative mating where animals prefer the somewhat similar to the exactly similar. Optimal outbreeding strikes a balance between inbreeding, which tends to preserve useful schemata against disruption through crossover, and outbreeding, which allows the incorporation of useful new mutations and schemata.

Schaffer & Eshelman (1991) suggested that genetic operators such as crossover could be characterized by a 'safety ratio', representing the ratio of the probability that an offspring will be better than its parents to the probability that it will be worse:  $P(\text{better})/P(\text{worse})$ . Optimal outbreeding can be viewed as a way of improving the safety ratio of crossover. It allows individuals to avoid breeding with others who are so similar as to be make recombination useless, and to avoid breeding with others so different as to make recombination too risky.

Eshelman & Schaffer (1991) developed an 'incest avoidance' method of preventing premature convergence by allowing individuals to mate only if their Hamming distance is above a certain threshold. This threshold drops according to some schedule as the population converges. They found that incest-prevention significantly increases the speed and robustness of finding global optima on 11 of 13 test functions. Allowing optimal outbreeding preferences to evolve can be seen as a way of automating this incest avoidance system. In Todd and Miller (1991) we developed a single-parameter method for specifying degree of desired outbreeding, and we found that even with a moderate degree of incest avoidance, speciation can still occur through assortative mating. Thus, the benefits of incest avoidance at the local level need not undermine the benefits of assortative mating at the global level.

The advantage of allowing outbreeding parameters to evolve becomes clear through the following argument. The key problem in crossover is to match the genetic correlation between parents to the 'correlation length' of the fitness landscape,

(Manderick et al., 1991). Longer correlation lengths allow parents with larger Hamming distances between them to still produce reasonably fit offspring. For rugged fitness landscapes with short correlation lengths, crossover works best between similar parents, but for smoother landscapes with longer correlation lengths, crossover between more dissimilar parents explores the space more efficiently (Kauffman, 1993). Since we may not know the correlation length of a fitness landscape ahead of time, we should allow outbreeding parameters to evolve as mate preferences.

## 6 Selective Mating I: Optimization Through Fitness-Based Sexual Selection

In selective mating, animals prefer mates that display some desired trait, regardless of their similarity on that trait. For example, peahens prefer peacocks with large tails, regardless of their own tail size. Whereas assortative mating affects mainly the genetic linkages among alleles, selective mating affects allele frequencies themselves, because animals with popularly-sought traits have more offspring than those with rarely-sought traits. So, whereas assortative mating tends to preserve genetic diversity, selective mating usually imposes directional or stabilizing selection that reduces diversity.

Why would animals ever use selective mating if natural selection is already in force? Natural selection is a powerful force over the long term, but it can be horribly noisy, irregular, inaccurate, and inefficient within each generation. In the wild, small fitness differences are often too weak to drive evolution very efficiently, and large fitness differences are often due to chance (e.g. chance effects in predation, competition, and pathogen exposure). Both problems reduce heritable variation in reproductive success, which is what drives evolution. Evolutionary computation avoids the first problem by using perfectly accurate, deterministic fitness functions, and avoids the second problem through various forms of fitness scaling or rank-based reproduction. In nature, selective mating helps overcome both problems.

Noise can arise at any point in the mapping from genotype to fitness. Noisy development can perturb the mapping from genotype to phenotype, and noisy evaluations can perturb the mapping from phenotype to fitness, either because phenotypes face a limited sample of the full range of performance tests during their lifetime, or because the performance tests themselves are noisy. In the wild, natural selection has no incentive to reduce this noise or to provide an evolving species with accurate information about the local fitness landscape. By contrast, selective mate preferences evolve precisely in order to reduce this noise and to provide individuals with maximal information about the genetic quality of potential mates. Selective mating can thereby overcome noise in several ways. First, if natural selection is directional but noisy, mate preferences can internalize the current average selection 'vector', thereby smoothing out and speeding up the evolutionary trajectory of the population. For example, if a fitness function tends to select for a particular trait, but with high levels of noise, mate preferences should evolve to favor that trait, and can impose sexual

selection with much less noise. Second, mate preferences could specify the sample size of performance evaluations to apply to each individual (just as female animals demand that male courtship displays be repeated a certain number of times before they choose a mate), so mate preferences could adjust the amount of noise across evaluations; this would be an automatic solution to the 'sample allocation problem' studied by Aizawa and Wah (1993) for noisy fitness functions. Third, mate preferences could also implement a form of adaptive sampling of the problem space by biasing search towards currently tricky problems (just as female animals prefer males who deliver difficult-to-find prey as nuptial gifts). Schultz (1991) found that this sort of difficulty-biased sampling leads to solutions more efficiently and robustly than uniform sampling.

Aside from the problem of noise, there is the problem of maintaining meaningful variation in reproductive success: fitness scores may be so similar that evolution happens very slowly. GA methods of fitness scaling, rank-based selection, and tournament competition have been developed to address this problem, but such methods require experimenters to set some scaling factor or bias term based on assumptions about the fitness landscape's structure. When fitness functions are noisy, such assumptions are especially tricky for rank-based and tournament methods (Fogarty, 1993). In nature, sexual selection is the fitness scaler: females convert slight differences in male quality into large differences in male reproductive success by mating disproportionately with higher-quality males. The scaling factor is determined by the degree of polygyny, which in turn is determined by evolved mate preferences. Each individual has strong incentives for using mate preferences that register slight differences in the fitness of potential mates, so as to give their offspring a slight adaptive advantage in the next generation. These preferences can maintain substantial variance in reproductive success even when natural-selective differences in viability are minimal. Thus, selective mate preferences tend to impose an automatic, flexible form of fitness scaling that evolves to match the fitness landscape's structure.

A final advantage of selective mating is that 'directional preferences' can implement a form of 'evolutionary momentum' (Miller & Todd, 1993), which is analogous to 'momentum adaptation' developed by Ostermeier (1992) for evolution strategies, or the 'momentum term' used in neural network back-propagation learning. That is, directional preferences tend to keep a population moving in a particular direction in phenotype space. Mate preferences should automatically evolve to exploit and internalize current gradient information about the local natural-selection vector. Momentum is most useful in smooth fitness landscapes, and in non-stationary landscapes where fitness peaks are moving around in a somewhat predictable way (e.g. under co-evolution). In this latter case, directional mate preferences could evolve to point in the direction of consistent peak-movement, to anticipate selection pressures generation-by-generation (see Todd & Miller, 1993).

## 7 Selective Mating II: Search Through Neophilic Sexual Selection

Aside from preferring similarity and quality in mates, many animals prefer novelty. The 'rare-male' effect results from female choice in favor of rare phenotypes, and is a powerful form of negative frequency dependent selection. Darwin (1871) took neophilia (preference for novelty) seriously as a force in sexual selection, implying that it could often lead to evolutionary innovations unrelated to ecological demands. Neophilia can be implemented in GAs by allowing individuals to sample several individuals and to favor the most unusual one; neophilia can be combined with other forms of assortative and selective mating, either as a weighted factor applied to the same phenotypic traits, or as a force applied to independent traits.

Without neophilia, useful new mutations and schemata have trouble establishing a toehold in an evolving population: because they start out at such a low frequency, their disappearance through genetic drift is often more likely than their propagation through positive selection. With neophilic sexual selection, new mutations and schemata get an immediate boost that protects them from extinction and promotes their spread. From an individual animal's perspective, any novelty in a sexually-mature potential mate is at least promising, because it has survived at least one round of natural selection. Of course, deleterious novelties will be selected out eventually. Whereas the mutation rate scales the number of raw innovations introduced into a population per generation, neophilia scales the number of potentially useful innovations propagated towards higher frequency in the population (see Miller & Todd, in press, for more detailed discussion). Variable mutation rate schedules are useful in genetic search (Fogarty, 1989), and variable amounts of neophilia could approximate such schedules by adjusting the tendency of populations to propagate new mutations.

Neophilic selective mating has surprisingly important evolutionary implications. One central problem in evolution is how to start with a population already converged to a (locally optimal) fitness peak, and move somewhere else that proves better (Harvey, 1992). Sewall Wright (1932) suggested "The problem of evolution as I see it is that of a mechanism by which the species may continually find its way from lower to higher peaks ... In order that this may occur, there must be some trial and error mechanism on a grand scale", i.e. a scale over and above that of mutation plus selection. Wright's shifting balance theory, based on the geographic isolation of demes, was one way to promote divergence away from locally optimal peaks. Each deme can undergo a slightly different 'adaptive walk' in search of a better fitness peak (Kauffman, 1993). But neophilic sexual selection can provoke much more powerful, directional adaptive walks than Wright's model allows.

Directional mate preferences implement an indirect form of neophilia because the directional preferences usually favor regions of phenotype space that are sparsely populated. In our investigations of directional selection even without explicit neophilia (Miller & Todd, 1993), we found that directional preferences are powerful ways of escaping from local fitness peaks, to allow an entire population to explore the surrounding phenotype space in a coherent, directional way. Directional preferences

tend to take populations on adaptive walks that have long 'foray lengths', capable of escaping from local optima in rugged fitness landscapes.

This sort of capricious exploration can be viewed as a form of 'runaway sexual selection', first postulated by Fisher (1930) and recently simulated by Collins and Jefferson (1992). Runaway is one of the most powerful evolutionary forces, accounting for some of the most important innovations seen in biology, such as the human brain (Miller, 1993). By contrast, mutation by itself simply creates 'mutational clouds' or 'quasi-species' whose density falls off exponentially with Hamming distance from fitness peaks (see Eigen, 1992); these mutational clouds are extremely weak at escaping from fitness peaks. 'Mutational search' relies on passive, diffuse, undirected replication errors to move populations away from local optima, whereas sexual selection relies on the active, focused, directed power of mate choice. The fact that most of the minor evolutionary innovations that distinguish species from one another are sexually-selected traits, suggests that most adaptive walks away from local optima occur under the power of sexual selection (see Miller & Todd, in press). In summary, neophilic sexual selection may be to macroevolution what mutation is to microevolution: the main source of potentially adaptive heritable variation.

## 8 Selective Mating III: Search Through Aesthetic Sexual Selection

Many design problems include both functional criteria and aesthetic criteria. The functional criteria are often easier to implement as an explicit fitness function. To evolve designs through aesthetic criteria, several researchers have fallen back on using human judges to award fitness scores to individuals. For example, Dawkins' (1986) Blind Watchmaker program allows human users to evolve 'biomorphs' on the screen through an evolutionary process. Caldwell and Johnson (1991) replaced the police sketch artist with a GA, by using human witnesses to judge the similarity of evolved faces to those of crime perpetrators, Sims (1992) has had spectacular success in evolving computer graphics images and dynamical systems based on human aesthetic judgments; Todd and Latham (1992) have also developed amazing computer graphics images through simulated evolution with human aesthetic input.

Humans can get tired of judging things during long GA runs, however. But how can aesthetic decisions be internalized into a GA selection system, when aesthetic criteria are so hard to formalize? Mate preferences offer one possibility. Consider a GA evolving two populations simultaneously: a population of designs subject to human aesthetic choice, and a population of 'preferences' (e.g. neural networks whose inputs are designs and whose outputs are aesthetic scores) that receive fitness scores for their ability to predict and match the human aesthetic choice. If the preferences are sufficiently rich and flexible, they should evolve to internalize the human user's aesthetic criteria; they could then be locked in and used as the fitness function that selects further designs, as a kind of aesthetic sexual selection. Such methods could prove useful in art, advertising, product design, and architecture; perhaps CAD

systems of the future will learn the user's aesthetic preferences through the simulated evolution of mate preferences, which can then guide the evolutionary generation of new designs.

Aesthetic criteria are useful for more than just evolving computer artwork. Many complex design problems are made more tractable by reliance on aesthetic heuristics such as symmetry, simplicity, streamlining, smoothness, and modularity. If these criteria were used even for purely functional design problems, greater search efficiency might result. Such criteria could either be internalized in mate preferences by evolving preferences to imitate human aesthetic decisions, or could evolve autonomously because they improve the functional efficiency of the designs they favor. The reasoning is that mate preferences can pick out phenotypic traits that are different from those on which natural selection acts, but that are highly correlated with natural-selective fitness. For example, the travelling salesman problem favors smooth, non-overlapping tours, but the explicit fitness function ('minimize total distance travelled') cannot directly select for smoothness or lack of overlap; these 'aesthetic' but heuristically useful criteria could evolve as selective mate preferences that guide evolution.

## 9 Conclusions and Prospects

Kauffman (1993) and others have suggested that populations tend to tune their mutation rates, crossover rates, and developmental mechanisms (e.g. degrees of epistasis) to keep themselves perched on the 'edge of chaos' such that they optimize evolution's ability to search complex fitness landscapes. This paper has suggested that evolvable mate preferences offer an even more powerful, direct, and biologically well-established way for populations to guide their evolutionary dynamics. Indeed, mechanisms for assortative and selective mating evolve precisely because they improve the way in which sexual recombination operates. If evolutionary computation methods fail to incorporate mate choice mechanisms into the entities under selection, they are unlikely to be exploiting the full power of recombination.

Sexual selection through mate choice was overlooked in biology for over a hundred years following Darwin (1871). The evolutionary computation community need not repeat the same mistake. Sexual selection is a powerful way of doing several things that natural selection finds difficult, such as maintaining genetic diversity, guiding sexual recombination, optimizing adaptations under weak or noisy natural selection, discovering and propagating evolutionary innovations, and implementing selection based on human aesthetic criteria. Natural and sexual selection are complementary because natural selection typically favors convergent evolution onto a few (locally) optimal solutions, whereas sexual selection often results in unpredictable, divergent patterns of evolution, with lineages speciating spontaneously and exploring the space of phenotypic possibilities according to their evolved mate preferences. As a method of combining exploitation and exploration, natural plus sexual selection is

much more powerful than selection plus mutation, because sexual selection delivers more complex, more integrated, more promising innovations than raw mutation does. In conclusion, the rich biological literature on sexual selection (see Andersson, 1994; Bateson, 1983; Cronin, 1991; Ridley, 1993) sets the stage for a sexual revolution in evolutionary computation.

## Acknowledgments

This research was supported by NSF-NATO Post-Doctoral Research Fellowship RCD-9255323 and NSF Research Grant INT-9203229. Thanks to Peter Todd and Inman Harvey for collaborating on the research, and to Dave Cliff and Phil Husbands for useful feedback.

## References

- Aizawa, A. N., & Wah, B. W. (1993). Dynamic control of genetic algorithms in a noisy environment. In Forrest (Ed.), pp. 48-55.
- Andersson, M. (1994). *Sexual selection*. Princeton U. Press.
- Bateson, P. (Ed.). (1983). *Mate choice*. Cambridge U. Press.
- Belew, R. K., & Booker, L. B. (Eds.). (1991). *Proc. of the Fourth Int'l Conf. on Genetic Algorithms*. Morgan Kaufmann.
- Booker, L.B. (1985). Improving the performance of genetic algorithms in classifier systems. In *Proc. of an Int'l Conf. on Genetic Algorithms and their Applications*, pp. 80-92.
- Caldwell, C., & V. S. Johnston (1991). Tracking a criminal suspect through "face-space" with a genetic algorithm. In Belew & Booker (Eds.), pp. 416-421.
- Cavicchio, D. J. (1970). *Adaptive search using simulated evolution*. Ph.D. thesis, Univ. Michigan.
- Collins, R. J., & Jefferson, D. R. (1991). Selection in massively parallel genetic algorithms. In Belew & Booker (Eds.), pp. 249-256.
- Collins, R. J., & Jefferson, D. R. (1992). The evolution of sexual selection and female choice. In Varela & Bourgoin (Eds.), pp. 327-336.
- Cronin, H. (1991). *The ant and the peacock: Altruism and sexual selection from Darwin to today*. Cambridge U. Press.
- Darwin, C. (1862). *On the various contrivances by which orchids are fertilized by insects*. John Murray.
- Darwin, C. (1871). *The descent of man, and selection in relation to sex (2 vols.)* John Murray.
- Davidor, Y., Yamada, T., & Nakano, R. (1993). The ECOlogical framework II: Improving GA performance at virtually zero cost. In Forrest (Ed.), pp. 171-176.
- Dawkins, R. (1986). *The blind watchmaker*. Norton.
- De Jong, K.A. (1975) *An analysis of the behavior of a class of genetic adaptive systems*. Ph.D. thesis, Univ. Michigan.
- Deb, K., & Goldberg, D.E. (1989) An investigation of niche and species formation in genetic function optimization. In Schaffer (Ed.), pp. 42-50.
- Dobzhansky, T. (1937). *Genetics and the origin of species*. (Reprinted 1982). Columbia U. Press.

- Eigen, M. (1992). *Steps towards life: A perspective on evolution*. Oxford U. Press.
- Eshelman, L. J., & Schaffer, J. D. (1991). Preventing premature convergence in genetic algorithms by preventing incest. In Belew & Booker (Eds.), pp. 115-122.
- Eshelman, L. J., & Schaffer, J. D. (1993). Crossover's niche. In Forrest (Ed.), pp. 9-14.
- Fisher, R. A. (1930). *The genetical theory of natural selection*. Clarendon Press.
- Fogarty, T. C. (1989). Varying the probability of mutation in the genetic algorithm. In Schaffer (Ed.), pp. 104-109.
- Fogarty, T. C. (1993). Reproduction, ranking, replacement, and noisy evaluations: Experimental results. In Forrest (Ed.), p. 634.
- Forrest, S. (Ed.). (1993). *Proc. of the Fifth Int'l Conf. on Genetic Algorithms*, Morgan Kaufmann.
- Goldberg, D.E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley.
- Goldberg, D.E., Deb, K., & Horn, J. (1992). Massive multimodality, deception, and genetic algorithms. In Männer & Manderick (Eds.), pp. 37-46.
- Goldberg, D.E., & Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization. In *Proc. of the Second Int'l Conf. on Genetic Algorithms*, pp. 41-49.
- Gordon, V. S., & Whitney, D. (1993). Serial and parallel genetic algorithms as function optimizers. In Forrest (Ed.), pp. 177-183.
- Gorges-Schleuter, M. (1989). ASPARAGOS: An asynchronous parallel genetic optimization strategy. In Schaffer (Ed.), pp. 422-427.
- Gorges-Schleuter, M. (1992). Comparison of local mating strategies in massively parallel genetic algorithms. In Männer & Manderick (Eds.), pp. 553-562.
- Harvey, I. (1992). Species adaptation genetic algorithms: The basis for a continuing SAGA. In Varela & Bourguine (Eds.), pp. 346-354.
- Hillis, W. D. (1992). Co-evolving parasites improve simulated evolution as an optimization procedure. In C. G. Langton et al. (Eds.), *Artificial Life II*, pp. 313-324. Addison-Wesley.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford U. Press.
- Kirkpatrick, M. (1987). The evolutionary forces acting on female preferences in polygynous animals. In Bradbury & Andersson (Eds.), pp. 67-82.
- Littman, M. J., & Ackley, D. H. (1991). Adaptation to constant utility non-stationary environments. In Belew & Booker (Eds.), pp. 136-142.
- Mahfoud, S. W. (1992). Crowding and preselection revisited. In Männer & Manderick (Eds.), pp. 27-36.
- Manderick, B., & Spiessens, P. (1989). Fine-grained parallel genetic algorithms. In Schaffer (Ed.), pp. 428-433.
- Manderick, B., de Weger, M., & Spiessens, P. (1991). The genetic algorithm and the structure of the fitness landscape. In Belew & Booker (Eds.), pp. 143-150.
- Männer, R., & Manderick, B. (Eds.). (1992). *Parallel problem solving from nature, 2*. North-Holland.
- Mayr, E. (1942). *Systematics and the origin of species*. (Reprinted 1982). Columbia U. Press.
- Michod, R. E., & Levin, B. R. (Eds.). (1988). *The evolution of sex: An examination of current ideas*. Sinauer.
- Miller, G. F. (1993). *Evolution of the human brain through runaway sexual selection*. Ph.D. thesis, Stanford University Psychology Department. (To be published as a book by MIT Press/Bradford Books, 1995).
- Miller, G. F. (in press). Sexual selection in human evolution: Review and prospects. For C. Crawford (Ed.), *Evolution and human behavior: Ideas, issues, and applications*. Lawrence Erlbaum.
- Miller, G. F., & Todd, P. M. (1993). Evolutionary wanderlust: Sexual selection with directional mate preferences. In J.-A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), *From Animals to Animals 2: Proc. Second Int'l Conf. on Simulation of Adaptive Behavior*, pp. 21-30. MIT Press.
- Miller, G. F., & Todd, P. M. (in press). The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. For W. Banzaf & F. Eckman (Eds.), *Proceedings of the 1992 Monterey Biocomputation Workshop*.
- Mühlenbein, H. (1989) Parallel genetic algorithms, population genetics, and combinatorial optimization. In Schaffer (Ed.), pp. 416-421.
- Mühlenbein, H. (1992). Darwin's continent cycle theory and its simulation by the Prisoner's Dilemma. In Varela & Bourguine (Eds.), pp. 236-244.
- Ostermeier, A. (1992). An evolution strategy with momentum adaptation of the random number distribution. In Männer & Manderick (Eds.), pp. 197-206.
- Pomiankowski, A. (1988). The evolution of female mate preferences for male genetic quality. *Oxford Surveys in Evolutionary Biology*, 5, 136-184.
- Ridley, M. (1993). *The red queen: Sex and the evolution of human nature*. Viking.
- Schaffer, J. D. (Ed.). (1989). *Proc. of the Third Int'l Conf. on Genetic Algorithms*. Morgan Kaufmann.
- Schaffer, J. D., & Eshelman, L. J. (1991). On crossover as an evolutionarily viable strategy. In Belew & Booker (Eds.), pp. 61-68.
- Schultz, A. C. (1991). Adapting the evaluation space to improve global learning. In Belew & Booker (Eds.), pp. 158-164.
- Sims, K. (1992). Interactive evolution of dynamical systems. In Varela & Bourguine (Eds.), pp. 171-178.
- Todd, P. M., & Miller, G. F. (1991). On the sympatric origin of species: Mercurial mating in the Quicksilver Model. In Belew & Booker (Eds.), pp. 547-554.
- Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, 2(1): 5-47.
- Todd, S., & Latham, W. (1992). Artificial life or surreal art? In Varela & Bourguine (Eds.), pp. 504-513.
- Varela, F. J., & Bourguine, P. (Eds.). (1992). *Toward a practice of autonomous systems: Proceedings of the First European Conference on Artificial Life*. MIT Press.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proc. Sixth. Int'l Congr. Genetics*, 356-366.