Parental and Cyclic-Rate Mutation in Genetic Algorithms: An Initial Investigation*

Theodore P. Hoehn and Chrisila C. Pettey
Department of Computer Science, Box 48, Middle Tennessee State University, Murfreesboro, TN 37132
cshoehn@knuth.mtsu.edu, cscbp@mtsu.edu
(615)898-2416

Abstract

Genetic algorithms (GAs) have traditionally combined and confounded the effects of gametic and somatic mutation. In an attempt to disentangle these effects, a novel, modified GA is proposed in which mutation occurs not only following selection (i.e., in offspring)—as is the tradition—but also at the *parental* level, between evaluation and selection of population members. Two experiments were conducted that compared the relative performances of several mutational-rate variants of the proposed parental mutation GA with corresponding "standard" GAs matched with them in terms of overall probability of mutation during an iteration of the algorithm. Results indicate that inclusion of parental mutation within a GA does not in general adversely affect its performance. For both types of GA, a novel, cyclic-rate mutation variant outperformed the low and (very) high constant-rate variants. Results of a third experiment indicate that, for the most part, it is irrelevant at which algorithmic locus (parental vs. offspring) the cyclic-rate mutation occurs or whether there is also mutation at the other locus.

1 Introduction

The most typical and general form of (sequential) genetic algorithm (GA), represented by Goldberg's Simple GA (SGA) [4], comprises population initialization and evaluation phases, followed by iteration of selection, recombination (i.e., mating/crossover), mutation, and evaluation phases. Traditionally, mutation in GAs occurs following fitness-based selection [2, 7]. Population members at this stage of the algorithm represent offspring of the previous generation produced either by mating/crossover of parental pairs or by cloning of individual parents. The effects of this *offspring mutation* are immediately apparent (i.e., phenotypically expressed) in the fitnesses assigned during the subsequent evaluation phase of the algorithm. In biological genetics, however, the importance of mutation relates primarily to changes within an organism's germ cells [9]. Such changes occur *prior* to reproduction and are "invisible" with respect to that individual's phenotype (i.e., morphology and behavior in nature, evaluation outcome value in GAs). Only in an indirect manner, through the viability and fitness of resultant offspring, is the original, parent organism's fitness impacted.

To perhaps better model natural processes [5], the following novel, modified sequential GA, termed a "parental mutation GA" (see Figure 1), is proposed in which mutation occurs not only in offspring but also at the *parental* level, between evaluation and selection. Selection, as usual, is based on phenotype/fitness values determined and assigned during the preceding evaluation phase, but now the underlying genotype of an individual may have been altered through subsequent mutation. Thus, the effects of such mutation are not phenotypically expressed in the parents. Rather, they are "hidden" within the parental genotypes. Individuals preferred for selection on the basis of their original phenotype/fitness values may now actually contain genotypes corresponding to much lower (or possibly higher) phenotypes. Only in the succeeding generation are these altered genotypes expressed phenotypically, either directly through cloning of the parent or indirectly via the product of mating/crossover of a parental pair. Of course, these genotypes could yet

* *Note:* This paper is based on the first author's Master thesis, submitted in August 1998 to Middle Tennessee State University, Murfreesboro, TN.

```
GA:
    begin
        initialization;
        evaluation;
        while not done
            begin
                parental mutation;
                selection;
                recombination;
                (offspring)mutation;
                evaluation;
            end;
    end.
```

Figure 1: Parental mutation genetic algorithm

be altered by means of standard, offspring mutation prior to the individuals being evaluated and assigned fitnesses for future selection.

Within this proposed GA scheme, traditional mutation at the offspring level would represent (primarily) somatic (i.e., body cell) mutations. (Because the altered genotype is potentially passed on to the following generation, there is an element of gametic mutation as well.) Somatic mutations, such as may be induced by radiation exposure and give rise to cancer, directly affect an individual's phenotype/fitness in terms of the individual's likelihood to survive and reproduce [9]. Parental mutation, on the other hand, would represent (purely) gametic (i.e., germ cell) mutations. As noted, these mutations, such as those underlying the appearance of a birth defect caused by some environmental mutagen, affect not the phenotype/fitness of the individual but rather that of its offspring [9]. In that sense, GA mutation at the parental level could perhaps also be viewed as modeling—in an admittedly crude manner—the nature of recessive genetic diseases [9], with parents serving as carriers of the "recessive," mutated genes.

The present study represents an initial investigation of the effects of this "hidden," parental mutation on GA performance. The relationship between various rates of mutation in a GA, particularly a novel, cyclic pattern, and performance is also examined.

2 Experiment 1A. Comparisons with Standard GA

Four GAs, each of which relies on the SGA basic design and structures [4], were compared in this experiment

Table 1: Mutation Rates in Experiment 1A

GA	Parents	Offspring
NPLO	none	.001
LPLO	.001	.001
HPLO	.1	.001
CPLO	.001, .1	.001

(see Table 1). Consistent with the "conventional wisdom" related in [8], all were implemented, in C++, with a population size of 50, a one-point crossover rate of .60, and an offspring mutation rate of .001. The selection strategy was proportional selection (i.e., stochastic sampling with replacement) with an elitist policy implementation under which, if called for, the elite individual replaced a randomly chosen current population member.

A "standard" GA, or *No-Parental, Low-Offspring mutation (NPLO)*, was included to serve as a baseline comparison for GA performance. There were three parental mutation GAs. *Low-Parental, Low-Offspring mutation (LPLO)* incorporates parental mutation set at a constant low rate (.001), and *High-Parental, Low-Offspring mutation (HPLO)* at a constant (very) high rate (.1). In *Cyclic-Parental, Low-Offspring mutation (CPLO)*, parental mutation is set at a low base rate (.001). However, every tenth parental generation, starting with the first, this mutation changes to a high periodic rate (.1). LPLO and HPLO are intended as rough models of the respective chronic parental mutation rates in nature (with the overstated .1 rate in HPLO perhaps corresponding to the effects of some environmental risk factor, such as atmospheric ozone depletion, taken to the extreme). The CPLO parental mutation rate, on the other hand, is intended to reflect conditions of acute "catastrophic" levels of mutation such as may result from nuclear war(s). Of course, the applicability of parental mutation, and of these particular mutation rates, within the context of GA behavior is independent of these suggested extensions to natural counterparts.

The four GAs were compared in terms of their performances in optimizing De Jong's functions F1-F5. A test on each function comprised 100 runs of the GA with 100 generations per run. Outcome data for each GA, reported as an average across the 100 runs, were composed of the best phenotype for each generation/population and the population-average phenotype for each generation.

For each GA, the same set of 100 seeds for the pseudo-

random number generator was used for the 100 runs on a function. Additionally, in the interest of experimental control, a second, distinct generator (always using the same seed value) was employed within the LPLO, HPLO, and CPLO GAs solely to provide random numbers for use in determinations of parental mutation. Thus, not only the population initialization for any particular run, but also all subsequent probabilistic operations, with the exception of parental mutation, were identical across all four GAs. Code for the pseudorandom number generators was generalized from that of the SGA generator [4].

In terms of the average ultimate-best phenotype (i.e., the best phenotype ever produced over the course of the 100 generations) for each GA, CPLO ranked either first or (a close) second for all five functions. NPLO, representing the baseline condition for GA performance, was outranked by all three parental mutation GAs on functions F1, F2, and F5, and by all but HPLO on F3 and F4. Although the three parental mutation GAs, particularly the cyclic-rate variant (CPLO), performed well relative to the NPLO standard GA, it could be argued that NPLO does not in fact represent the appropriate baseline condition. This possibility is examined in Experiment 1B.

3 Experiment 1B. Matched-Low Standard GA as Baseline Condition

The three parental mutation GAs tested in Experiment 1A include parental mutation (at whichever rate) *in addition to* offspring mutation (at a .001 rate). Hence, the overall probability of mutation during an iteration of the GA is higher with any of them than it is with the NPLO standard GA. If it is this overall probability, rather than or along with the algorithmic locus of mutational activity (i.e., at the parental vs. offspring level), that is determinant of performance, then the comparison between NPLO and the parental mutation GAs in Experiment 1A does not constitute a fair test of the relative efficacy of the four GAs.

To address this concern, a different standard GA was devised for Experiment 1B to act as the baseline comparison of performance for the three parental mutation GAs (see Table 2). This new GA, *Matched-Low Offspring mutation (MLO)*, contains offspring mutation set at a rate that matches the cumulative probability of parental and offspring mutation in the constant low-rate parental mutation GA (LPLO). Except for the substitution of NPLO with MLO as the baseline condition GA, Experiment 1B was identical in all respects to Experiment 1A.

Table 2: Mutation Rates in Experiment 1B

GA	Parents	Offspring
MLO	none	.001998
LPLO	.001	.001
HPLO	.1	.001
CPLO	.001, .1	.001

The new baseline condition, MLO, was no longer the clearly worst performing GA in terms of average ultimate-best phenotype. Indeed, it was ranked last for none of the five functions. MLO outranked all three parental mutation GAs on F4, and was ranked second on F3 and third on functions F1, F2, and F5. The results of this experiment call into question the apparent superiority in performance of parental mutation GAs over standard GAs evidenced in Experiment 1A. Experiment 2 continues this line of investigation.

4 Experiment 2. Differential Rates of Mutation in Standard GA

When, in Experiment 1B, the overall probability of mutation during an iteration of the algorithm was equated between the constant low-rate parental mutation GA (LPLO) and the standard GA being used for baseline comparison (MLO in this case), the performances of these two GAs were found to be comparable. The constant high-rate parental mutation GA (HPLO) still outperformed the other three GAs on functions F2 and F5, and the cyclic-rate variant (CPLO) still excelled on F1 and F3 (and did quite well on the other functions also). The results of Experiment 1B regarding the LPLO-MLO comparison, however, make it an open question as to whether the superiority in performance with parental mutation GAs evidenced with LPLO, HPLO, and CPLO in Experiment 1A and with HPLO and CPLO in Experiment 1B are attributable to factors intrinsic to parental mutation GAs or to some other, extrinsic factor (i.e., the overall probability of mutation during an iteration of the GA).

Experiment 2 consequently was designed to assess the relative performances of standard GAs under conditions of constant low, constant high, or cyclic rates of (offspring) mutation. As with the Matched-Low Offspring mutation (MLO) GA utilized in Experiment 1B, each of the three standard GAs compared in this experiment (see Table 3) contains offspring mutation set at a rate that matches the cumulative probability of parental and offspring mutation in a corresponding

Table 3: Mutation Rates in Experiment 2

GA	Parents	Offspring
MLO	none	.001998
MHO	none	.1008
MCO	none	.001998, .1008

parental mutation GA. In addition to MLO (which corresponds to LPLO), the experiment compares the performances of a *Matched-High Offspring mutation (MHO)* GA, corresponding to HPLO, and a *Matched-Cyclic Offspring mutation (MCO)* GA, corresponding to CPLO. Except for the substitution of "matched standard" GAs for the corresponding parental mutation GAs (and the elimination of any GA added as an external baseline condition), Experiment 2 was identical in all respects to Experiments 1A-B.

In contrast to the case with Experiments 1A-B, outcome data of this experiment underwent more than visual inspection. For each function, results with respect to ultimate-best phenotype produced per run of a GA were submitted to statistical analysis (with a minimum two-tailed alpha level of .05 throughout). Because the form of the outcome distributions tended to be widely divergent between GAs, the Friedman test, a nonparametric analysis for two or more treatment levels [1], was utilized as the overall test of statistical significance for observed differences in performance on a function. This test views the experimental treatments as a within-subjects factor (i.e., as repeated measures). In this case, the pseudorandom number generator seeds represent individual subjects and the various GAs the treatments applied to each subject. That is, because the same set of 100 seeds was used in conducting the 100 runs for every GA, each seed could be considered as being one of 100 subjects administered every level of treatment. As a nonparametric measure, the unit of analysis for the test is the relative rank of each treatment outcome for a subject, rather than the actual numerical value of that outcome. If the Friedman test that was carried out on the data for a particular function indicated a *significant effect of the treatment factor on performance*, which means that results, for at least some of the treatment levels, differ significantly among themselves, pairwise multiple comparisons of the mean ranks for all the treatments (i.e., GAs) subsequently were performed to identify the relevant comparisons. This was accomplished using a test, analogous to Fisher's Least Significant Difference, based on the t statistic [1].

Table 4 presents the mean rank (across 100 runs) of the ultimate-best phenotype produced by each GA in Experiment 2, and its relative ranking, for each function. With regard to this data, the effect of GA mutational-rate variant on performance was significant at the $p < .001$ level in all five Friedman tests ($\chi^2(2,\ N = 100) = 20.615, 25.235, 75.125, 142.940$, and 55.580 for functions F1-F5, respectively). Consequently, for each function, all possible pairwise comparisons (t(198)) between mean ranks for the three GAs were analyzed. Results of these statistical analyses are indicated in the table.

For purposes of comparison, Table 5 presents the mean rank (across 100 runs) of the ultimate-best phenotype produced by each parental mutation GA tested in Experiments 1A-B, and its relative ranking, for each function. With regard to this data, the effect of GA mutational-rate variant on performance was significant at the $p < .001$ level in all five Friedman tests ($\chi^2(2,\ N = 100) = 38.385, 39.260, 57.165, 148.460$, and 62.000 for functions F1-F5, respectively). Consequently, for each function, all possible pairwise comparisons (t(198)) between mean ranks for the three GAs were analyzed. Results of these statistical analyses are indicated in the table.

The results of the statistical analyses reported above (and the data presented in Tables 4 and 5) obviously exhibit very similar patterns for the parental mutation GAs tested in Experiments 1A-B and the corresponding matched standard GAs tested in Experiment 2. In particular, the cyclic-rate mutation GA variants (CPLO and MCO, respectively) overall outperformed their two respective constant-rate variants in terms of ultimate-best phenotypes produced. Thus, the three different rates of mutation (constant low, constant high, or cyclic) used in Experiments 1A-B and 2 seem to have analogous effects on performance whether the mutation is incorporated within a standard GA or a parental mutation GA. However, even though the *relative* performances of the three mutational rate variants are comparable across the two types of GA, there could still be significant differences in the outcome values for any particular variant between the parental mutation and matched standard GAs. That is to say, the parental mutation GAs might outperform their standard GA counterparts (or vice versa) yet still give rise to a similar pattern of results across mutational rate variants. The following tables, analyses, and associated discussion examine this possibility in detail.

Table 6 presents the mean rank (across 100 runs) of the ultimate-best phenotype produced by each paired parental mutation and matched standard GA, and its

Table 4: Mean Rank of Ultimate-Best Phenotypes (and Relative Rankings) for Matched Standard GAs in Experiment 2 ($*\vee**p < .05$ or $p < .01$ for all pairwise comparisons in column. $**p < .01$ for all pairwise comparisons in column except as indicated by bracketing, | |, of a nonsignificant comparison.)

GA	F1	F2	F3	F4	F5
	**	**	**	$*\vee$**	**
MLO	|2.09 (2)|	2.38 (3)	1.98 (2)	1.42 (1)	2.58 (3)
MHO	|2.27 (3)|	|1.68 (1)|	2.63 (3)	2.97 (3)	1.55 (1)
MCO	1.64 (1)	|1.93 (2)|	1.40 (1)	1.61 (2)	1.87 (2)

Table 5: Mean Rank of Ultimate-Best Phenotypes (and Relative Rankings) for Parental Mutation GAs from Experiments 1A-B ($*\vee**p < .05$ or $p < .01$ for all pairwise comparisons in column. $**p < .01$ for all pairwise comparisons in column.)

GA	F1	F2	F3	F4	F5
	**	**	**	$*\vee$**	**
LPLO	1.94 (2)	2.47 (3)	1.95 (2)	1.42 (1)	2.60 (3)
HPLO	2.46 (3)	1.94 (2)	2.56 (3)	2.99 (3)	1.50 (1)
CPLO	1.60 (1)	1.59 (1)	1.50 (1)	1.59 (2)	1.90 (2)

relative ranking within the pair, for each function. With regard to this data, the effect of GA type (i.e., parental mutation or matched standard) was not significant in any of the Friedman tests comparing LPLO and MLO performances for functions F1-F5. However, for every function except F3, the effect of GA type was significant in the Friedman tests comparing HPLO and MHO performances ($\chi^2(1, N = 100) = 8.41$, 7.84, and 9.00, $p < .01$, for F1, F2, and F4, respectively; $\chi^2(1, N = 100) = 5.29$, $p < .05$, for F5). Finally, only for functions F2 and F5 ($\chi^2(1, N = 100) = 4.41$ and 5.29, respectively) was the effect of GA type significant ($p < .05$) in the Friedman tests comparing CPLO and MCO performances.

The overall probability of mutation (.001998) differentially partitioned within the parental mutation GA LPLO and its matched standard GA MLO is very small. Hence, the lack of significant differences in performance between these two GAs on any of De Jong's functions is perhaps not particularly surprising in any case. The overall probability of mutation (.1008) differentially partitioned within the parental mutation GA HPLO and its matched standard GA MHO, in contrast, is very high. It might well be expected that any intrinsic differences in efficacy between the two types of GA under comparison would be revealed under such extreme conditions. This expectation was borne out in that MHO significantly outperformed HPLO on all functions except F3. This finding indicates that too much parental mutation adversely affects performance relative to that same overall rate applied solely at the offspring level.

The parental mutation GA-matched standard GA comparison that is most interesting, based on the results of Experiments 1A-B and 2 comparing the relative performances of the various mutational rate variants, is that of the two respective cyclic-rate mutation GAs. These GAs performed equivalently on functions F1, F3, and F4. However, the parental mutation GA (CPLO) significantly outperformed its matched standard GA (MCO) on function F2, whereas the opposite outcome was true for F5. This mixed pattern of results across functions is not easily interpreted, but will be explored further in Experiment 3.

5 Experiment 3. Cyclic-Rate Mutation Across Algorithmic Loci

The novel, cyclic pattern was found to be the overall best performing rate of mutation in the GAs compared across Experiments 1A-B and 2. Experiment 3 was designed to examine the nature of this effect in more depth. It assesses the relative performances of four cyclic-rate mutation GAs (see Table 7), each representing a variant based on the algorithmic locus (i.e., parental or offspring) at which the cyclic-

Table 6: Mean Rank of Ultimate-Best Phenotypes (and Relative Rankings) for Parental Mutation and Matched Standard GAs Across Experiments 1A-B and 2 (*$p < .05$. **$p < .01$.)

GA	F1	F2	F3	F4	F5
LPLO	1.48 (1)	1.55 (2)	1.52 (2)	1.56 (2)	1.54 (2)
MLO	1.52 (2)	1.45 (1)	1.48 (1)	1.44 (1)	1.46 (1)
	**	**		**	*
HPLO	1.64 (2)	1.64 (2)	1.53 (2)	1.65 (2)	1.62 (2)
MHO	1.36 (1)	1.36 (1)	1.47 (1)	1.35 (1)	1.38 (1)
		*			*
CPLO	1.49 (1)	1.39 (1)	1.55 (2)	1.44 (1)	1.62 (2)
MCO	1.51 (2)	1.61 (2)	1.45 (1)	1.56 (2)	1.38 (1)

rate mutation occurs and whether there is, in addition to the aforementioned mutation, either constant low-rate (.001) or no mutation at the other locus. As with the parental mutation and corresponding matched standard GAs from the previous experiments, all of the cyclic-rate mutation GAs compared in this experiment were equated for their overall probability of mutation during an iteration of the GA. In addition to the Matched-Cyclic Offspring mutation (MCO) GA tested in Experiment 2 and the Cyclic-Parental, Low-Offspring mutation (CPLO) GA tested in Experiments 1A-B, two new cyclic-rate mutation GA variants were devised: *Low-Parental, Cyclic-Offspring mutation (LPCO)* and *Matched-Cyclic Parental mutation (MCP)*. MCO and LPCO are identical, each involving cyclic-rate mutation at the offspring level, except that the latter GA includes (constant low-rate) mutation at the parental level as well. Similarly, MCP and CPLO are identical, each involving cyclic-rate mutation at the parental level, except that the latter GA includes (constant low-rate) mutation at the offspring level as well. From another perspective, MCO and MCP can be viewed as similar in that both involve (cyclic-rate) mutation at only a single algorithmic locus (offspring and parental, respectively), whereas LPCO and CPLO involve (cyclic-rate plus constant low-rate) mutation at both algorithmic loci. It must be stressed that for all four GAs the overall probability of mutation during an iteration of the algorithm is the same, and is cyclic in nature. Only the distribution of mutational opportunities across algorithmic loci differs between GAs.

Except for the substitution of cyclic-rate mutation GAs for parental mutation or matched standard GAs, Experiment 3 was identical in all respects to Experiments 1A-B and 2. Statistical analyses were identical to those discussed above in reference to Experiment 2 and consisted of initial Friedman tests, one per function, followed by multiple comparisons within indicated functions.

Table 7: Mutation Rates in Experiment 3

GA	Parents	Offspring
MCO	none	.001998, .1008
LPCO	.001	.001, .1
CPLO	.001, .1	.001
MCP	001998, .1008	none

Table 8 presents the mean rank (across 100 runs) of the ultimate-best phenotype produced by each GA, and its relative ranking, for each function. With regard to this data, the effect of cyclic-rate mutation GA variant on performance was significant in the Friedman tests only for functions F2 ($\chi^2(3, N = 100) = 9.375, p < .05$) and F5 ($\chi^2(3, N = 100) = 23.457, p < .001$). Consequently, for each of these two functions, all possible pairwise comparisons (t(297)) between mean ranks for the four GAs were analyzed. Results of these statistical analyses are indicated in the table.

The results of the present experiment extend the findings regarding the cyclic-rate parental mutation GA vs. matched standard GA comparisons reported in Experiment 2. In both sets of experiments, the cyclic-rate mutation GA variants (CPLO and MCO across Experiments 1A-B and 2; MCO, LPCO, CPLO, and MCP in this experiment) performed equivalently on functions F1, F3, and F4. The significant differences in performance on F2 and F5 noted in Experiment 2, furthermore, are here observed on a finer scale. Across Experiments 1A-B and 2, CPLO outperformed MCO on F2; in this experiment, CPLO outperformed all of the other three variants, which were comparable

Table 8: Mean Rank of Ultimate-Best Phenotypes (and Relative Rankings) for Cyclic-Rate Mutation GAs in Experiment 3 ($*\vee**p < .05$ or $p < .01$ for all pairwise comparisons in column except as indicated by bracketing, | |, of nonsignificant comparisons. $**p < .01$ for all pairwise comparisons in column except as indicated by bracketing, | | or || ||, of a nonsignificant comparison).

GA	F1	F2	F3	F4	F5
		$*\vee**$			**
MCO	2.47 (2)	|2.56 (2)|	2.50 (3)	2.57 (4)	|2.28 (2)|
LPCO	2.33 (1)	|2.67 (4)|	2.37 (1)	2.48 (2)	|2.12 (1)|
CPLO	2.49 (3)	2.16 (1)	2.68 (4)	2.39 (1)	||2.75 (3)||
MCP	2.70 (4)	|2.61 (3)|	2.45 (2)	2.56 (3)	||2.86 (4)||

in performance among themselves. Thus, this particular cyclic-rate mutation GA variant, representing a cyclic-rate parental mutation/low constant-rate offspring mutation combination, seems to excel on this particular function. Across Experiments 1A-B and 2, MCO outperformed CPLO on F5. In this experiment, on F5 the two cyclic-rate mutation GA variants in which the cyclic rate applied to the offspring level (MCO and LPCO) outperformed the two variants in which the cyclic rate applied to the parental level (CPLO and MCP); the presence or absence of low constant-rate mutation at the other algorithmic locus had no significant effect on performance. Again, this mixed pattern of results across functions is difficult to interpret, although in general, these four GAs largely performed alike.

6 Conclusions and Future Research

The results of this study clearly indicate that inclusion of "hidden," parental mutation within a GA does not in general adversely affect its performance. Initial performance comparisons between various parental mutation GAs and a standard GA in Experiment 1A gave the impression that the parental mutation GAs actually outperformed the standard type (see also [6]). Subsequent comparisons in Experiments 1B and 2 between parental mutation and standard GAs equated for their overall probability of mutation during an iteration of the algorithm, however, did not support this initial impression. Nonetheless, the parental mutation GAs, if not superior, were not generally inferior to these standard GAs in performance.

Three different core rates of mutation were utilized in this study: constant low (.001), constant (very) high (.1), and cyclic (.001, .1). Those GAs, whether parental mutation or standard, that employed the cyclic pattern of mutation rate introduced here (and in [6]) overall outperformed their constant-rate peers. A cyclic-rate mutation GA may operate in important respects in a manner analogous to exponentially decreasing the (offspring) mutation rate over generations (e.g., [3]). That is to say, the periods of low mutation following each episode of high mutation may correspond in their effects to cyclic decreases in rate, during which intervals schema exploitation (i.e., detailed search of the local problem area) is emphasized over exploration (i.e., global search of the problem space). In any case, the principal distinguishing characteristic of a cyclic mutation rate, and the reason underlying its observed performance advantages over constant rates of mutation, would seem to be its periodic insertions of new genetic material into the population gene pool, thereby maintaining diversity.

The results of Experiment 3 assessing the relative performances of various cyclic-rate mutation GAs indicate that, for the most part, it is irrelevant at which algorithmic locus (parental vs. offspring) the cyclic-rate mutation occurs or whether there is also constant low-rate mutation at the other locus. There was, however, a somewhat mixed pattern of results regarding ultimate-best phenotypes for the four GAs evidenced across the five evaluation functions, suggesting some relationship between performance and the implementational factors differentiating these GAs.

In light of the present findings, several avenues for further research, regarding both parental and cyclic-rate mutation in GAs, suggest themselves:

- 1. The function-dependent results observed with cyclic-rate mutation GAs varying among themselves in terms of which algorithmic locus (parental vs. offspring) the cyclic-rate mutation occurs or whether there is also constant low-rate mutation at the other locus should be explored in more depth. For example, performance on evalu-

ation functions other than F1-F5 can be assessed.

- 2. Different values for the cyclic-rate mutation base and periodic settings could be utilized.
- 3. In a similar vein, the periodicity of the cyclic-rate mutation (every 10th generation in the present study) can be manipulated.
- 4. A mutation rate other than constant low could be applied at the algorithmic locus not subjected to cyclic-rate mutation.
- 5. Performances of the assorted parental mutation, matched standard, and cyclic-rate mutation GAs could be assessed and compared under a nonelitist selection policy (see [6]). There is some suspicion that although elitism may improve local search, it might do so at the expense of global coverage of the problem space [2].
- 6. As noted earlier in this paper, mutation at the offspring level in the parental mutation GAs as here constructed necessarily (as in standard GAs) reflects aspects of gametic mutation as well as somatic mutation. To remedy this, a GA could be devised that operates upon a population of bichromosomal individuals. One chromosome would represent somatic cells—for use in offspring mutation and in the evaluation and selection phases of the algorithm. The other chromosome would represent gametic cells—for use in parental mutation and in the recombination phase of the algorithm. This form of parental mutation GA would constitute a considerably greater departure from the standard sequential GA than does the form proposed and investigated in the present study. However, its more complete bifurcation of gametic and somatic mutation holds the promise of potentially interesting and rewarding results in terms of both naturalistic evolutionary modeling and practical GA applications.

References

[1] Conover, W.J., *Practical Nonparametric Statistics* (2nd ed.), John Wiley & Sons, New York, 1980.

[2] De Jong, K.A., "An Analysis of the Behavior of a Class of Genetic Adaptive Systems," (Doctoral dissertation, University of Michigan), Dissertation Abstracts International, 36, 5140B, (University Microfilms No. 76-9381), 1975.

[3] Fogarty, T.C., "Varying the Probability of Mutation in the Genetic Algorithm," *Proceedings of the Third International Conference on Genetic Algorithms*, J.D. Schaffer, ed., Morgan Kaufmann, San Mateo, CA, 1989.

[4] Goldberg, D.E., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA, 1989.

[5] Goldberg, D.E., "Zen and the Art of Genetic Algorithms," *Proceedings of the Third International Conference on Genetic Algorithms*, J.D. Schaffer, ed., Morgan Kaufmann, San Mateo, CA, 1989.

[6] Hoehn, T.P., "Wolves in Sheep's Clothing? The Effects of 'Hidden,' Parental Mutation on Genetic Algorithm Performance," *Proceedings of the 36th Annual ACM Southeast Conference*, K.N. King, ed., ACM Press, New York, 1998.

[7] Holland, J.H., *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, MIT Press, Cambridge, MA, 1992.

[8] Schaffer, J.D., R.A. Caruana, L.J. Eshelman, and R. Das, "A Study of Control Parameters Affecting Online Performance of Genetic Algorithms for Function Optimization," *Proceedings of the Third International Conference on Genetic Algorithms*, J.D. Schaffer, ed., Morgan Kaufmann, San Mateo, CA, 1989.

[9] Strickberger, M.W., *Genetics* (3rd ed.), Macmillan, New York, 1985.