

**Compressed Introns in a Linkage Learning  
Genetic Algorithm**

**Fernando G. Lobo  
Kalyanmoy Deb  
David E. Goldberg  
Georges R. Harik  
Liwei Wang**

IlliGAL Report No. 97010  
December 1997

Illinois Genetic Algorithms Laboratory  
Department of General Engineering  
University of Illinois at Urbana-Champaign  
117 Transportation Building  
104 South Mathews Avenue  
Urbana, IL 61801

# Compressed Introns in a Linkage Learning Genetic Algorithm

Fernando G. Lobo <sup>\*</sup>, Kalyanmoy Deb <sup>†</sup>, David E. Goldberg,  
Georges R. Harik, and Liwei Wang  
Illinois Genetic Algorithms Laboratory  
Department of General Engineering  
University of Illinois at Urbana-Champaign

## Abstract

Over the last 10 years, many efforts have been made to design a competent genetic algorithm. This paper revisits and extends the latest of such efforts—the linkage learning genetic algorithm. Specifically, it introduces an efficient mechanism for representing the non-coding material. Recent investigations suggest that this new method is crucial for solving a large class of hard optimization problems.

## 1 Introduction

The simple genetic algorithm (SGA) has been applied successfully in a variety of applications, including medical, financial, and all kinds of engineering problems. Its power comes from its ability to combine good pieces (building blocks) from different solutions and assemble them into a single super solution. But despite their success, there are still problems whose solution can be constructed by the juxtaposition of building blocks, and yet the SGA fails. The reason behind this failure is well understood. Over the last decade, Goldberg and his students have made many efforts to design a GA that can overcome this failure. The messy GA (Goldberg, Korb, & Deb, 1989), (Goldberg, Deb, Kargupta, & Harik, 1993), and more recently the gene expression messy GA (Kargupta, 1996) and the linkage learning GA (Harik, 1997) are examples of such efforts. The goal is to design a GA that can solve problems of bounded difficulty quickly, reliably, and accurately—a *competent GA*. Designing such an algorithm has been a difficult and challenging task, and the job is still partially complete. This paper extends the latest of these efforts, the linkage learning genetic algorithm (LLGA), with a technique that seems promising for solving a large class of hard optimization problems.

The next two sections provide background material needed to understand this paper. Section 2 explains the limitations of SGAs and section 3 reviews the linkage learning genetic algorithm. Then we describe a technique that doesn't change the behavior of the LLGA, but makes it computationally tractable for problems that require a large number of non-coding genes (more on this later). Finally, section 5 presents an empirical study of the time complexity of the LLGA on problems with exponentially scaled building blocks.

---

<sup>\*</sup>Fernando Lobo is a visiting student from “Grupo de Análise de Sistemas Ambientais”, FCT/UNL, Portugal.

<sup>†</sup>Kalyanmoy Deb is with the Department of Mechanical Engineering, IIT Kanpur, India.

## 2 The limitations of SGAs

Genetic algorithms are successful when they can propagate building blocks from generation to generation and combine them into a single solution. Recall that a building block (BB) is an association of features (bits in a binary coded GA) that as a whole give a high contribution to the fitness of an individual. The reason why SGAs fail is because they depend on the ordering that is imposed on the features of the problem. Within the SGA, each feature is coded at a fixed position in the chromosome. If the features that constitute a BB are located near each other, the SGA has no trouble propagating it. After all, Holland’s schema theory (Holland, 1975) says that short, highly fit similarities will tend to be propagated to subsequent generations. But what if the BB is not short? What if the BB is spread out along the chromosome? The answer is: SGAs fail because the BBs from the different individuals cannot be combined (mixed) in a reasonable amount of time. Thierens and Goldberg (1993) showed that the time needed to solve such problems grows exponentially as the problem size increases.

Ideally the BBs would be coded so that its genes are close together in the chromosome. But which genes constitute a BB is unknown beforehand. This is the great limitation of SGAs. A competent GA should be able to detect which genes constitute a BB and be able to propagate it as a whole chunk—as a super gene. A way to achieve this is to have a GA that autonomously rearranges the genes in the chromosome so that the genes that are related come closer in the chromosome—become tightly linked.

### 2.1 Test problems

In order to test new GA mechanisms, it is useful to create artificial problems with a crisp definition of building blocks—problems of bounded difficulty. These problems are constituted by a number of non-overlapping building blocks of a maximum size  $k$ . The overall fitness function is simply the sum of the fitnesses of the individual building blocks. These types of problems can be further divided into subclasses:

- Problems with uniformly scaled building blocks.
- Problems with exponentially scaled building blocks.
- A mixture of the above two types.

As the name suggests, uniformly scaled BBs are BBs that give the same contribution to the fitness function. Exponentially scaled means that the fitness of each BB is scaled by some constant power. Some can argue that these types of problems are not representative of the problems encountered in real world applications. This is partially true and it is very unlikely that a real world problem has such crisp definitions. Nonetheless, we argue that most problems have some kind of BB-structure. After all, if problems had no structure at all, then a genetic algorithm would be no better than random search. The success that GAs have had in numerous applications supports the hypothesis that problems do have a BB-like structure, possibly with some degree of interaction among BBs.

Our first goal is to design a GA that can efficiently solve problems with no interactions among BBs. Once we have such a GA, the next step will be to study the effects of having interactions among BBs. This last topic is sometimes called crosstalk. A competent GA should be able to solve any of these problems quickly, reliably, and accurately. Such a GA doesn’t fully exist yet, but we are closer than ever. Specifically, the linkage learning genetic algorithm (Harik, 1997) excels in problems with exponentially scaled BBs, but it is not as strong when the BBs are uniformly scaled. The next section reviews this algorithm.

### 3 The linkage learning genetic algorithm

This section gives a brief overview of the mechanics and operation of the linkage learning genetic algorithm (LLGA). For a more detailed description and analyses you should refer to Harik’s dissertation (Harik, 1997).

On problems with exponentially scaled building blocks, the LLGA autonomously rearranges the genes in the chromosome so that the genes that are related come closer in the chromosome—become tight. Once the linkage is tight, the LLGA has no trouble propagating the BBs. On these problems, the LLGA efficiently learns the linkage of each BB, one after the other. It first learns the linkage of the most important BB, and once this is formed, it moves on to next most important BB, and so on. It does so without ever losing diversity on the other genes due to a probabilistic expression mechanism. The remainder of this section describes the representation of the chromosome, the probabilistic expression mechanism, and the crossover operator.

#### 3.1 Representation

The LLGA represents each gene by a (locus,allele) pair. There is no under-specification, but over-specification is always present. Every gene has always both allele values represented in the chromosome. One of them is expressed and sent to the fitness function according to a probabilistic scheme. Having the complement allele for every gene ensures that diversity is never lost. In addition, a chromosome also contains non-coding genes (sometimes called introns in the evolutionary computation community). These genes give no contribution to the fitness function, and their purpose is to facilitate the propagation of BBs and the formation of linkage. Other researchers have also investigated the effects of including non-coding material in the chromosome structures (Levenick, 1991), (Levenick, 1995), (Wu & Lindsay, 1995). These studies are motivated by the existence of non-coding DNA in biological systems. For example, all living creatures contain large amounts of non-coding DNA. Specifically, about 97% of the human DNA is non-coding.

It is convenient to visualize the chromosome as a circular list of genes. Figure 1 gives an example of such a chromosome. The example is for a 3-bit problem. Genes 1,2, and 3 are the coding genes. Genes 4 and 5 are non-coding genes.




Figure 1: Example of a chromosome in the LLGA. Genes 1,2,3 represent features of the problem. Genes 4,5 are non-coding. Starting from the interpretation point and traversing the chromosome in a clockwise direction, the chromosome would be expressed as (3,1)(2,0)(1,1). Note the existence of additional genes with the complement alleles.

Each chromosome has an interpretation point. Starting from this point, the genes are expressed by traversing the chromosome circle in a clockwise manner and recording the first occurrence of each gene. In the example, we would get (3,1)(2,0)(1,1). These 3 genes would be sent to the fitness function. Genes number 4 and 5 are not sent to the fitness function because they are non-coding. Note that for every gene, there is always at least one copy with the complement allele. Harik named this expression mechanism as probabilistic expression because each chromosome can express a set of different solutions depending on the location of the interpretation point. What we have just described is the extended probabilistic expression (EPE-2). In this scheme, each gene has 1 copy of the expressed allele and at most 2 copies of the unexpressed allele (Harik, 1997). Next, we illustrate the crossover operator of the LLGA.

### 3.2 Crossover

The crossover operator works as follows. First, two chromosomes are randomly chosen from the population. One is named the recipient and the other is named the donor. A segment of the donor's chromosome is chosen randomly and is injected at a randomly chosen point (the grafting point) of the recipient. Then, the recipient chromosome is traversed in a clockwise manner and extra genes are deleted so that the resulting chromosome is a legal EPE-2 chromosome. Figure 2 illustrates the 3 steps of the crossover operator. When the selection rate is properly chosen, this crossover operator brings the genes that constitute a BB closer together. Details of why this happens are explained in Harik's dissertation. Once the linkage is tight, the BB has a very low chance of being disrupted and it is easily propagated to the other population members.

The generational loop of the LLGA is similar to the one of the SGA. Selection and crossover are applied generation after generation. Unlike the messy GA, there's no explicit primordial and juxtapositional phases. Harik showed that the LLGA efficiently solve a large class of hard problems that are difficult for the SGA—problems with exponentially scaled BBs. However, he also recognized that the LLGA didn't perform as well on uniformly scaled problems as it did on exponentially scaled ones. To attack this problem, a new interpretation scheme has been recently investigated (Wang, 1997). The main idea of this new interpretation scheme is to let the building blocks form is different parts of the chromosome circle. Preliminary results suggest that these kinds of problems can be solved efficiently, but in order to do so, the LLGA requires a very large number of non-coding genes. This number grows exponentially with the number of BBs. The next section introduces a technique that is capable of dealing with this exponential growth.

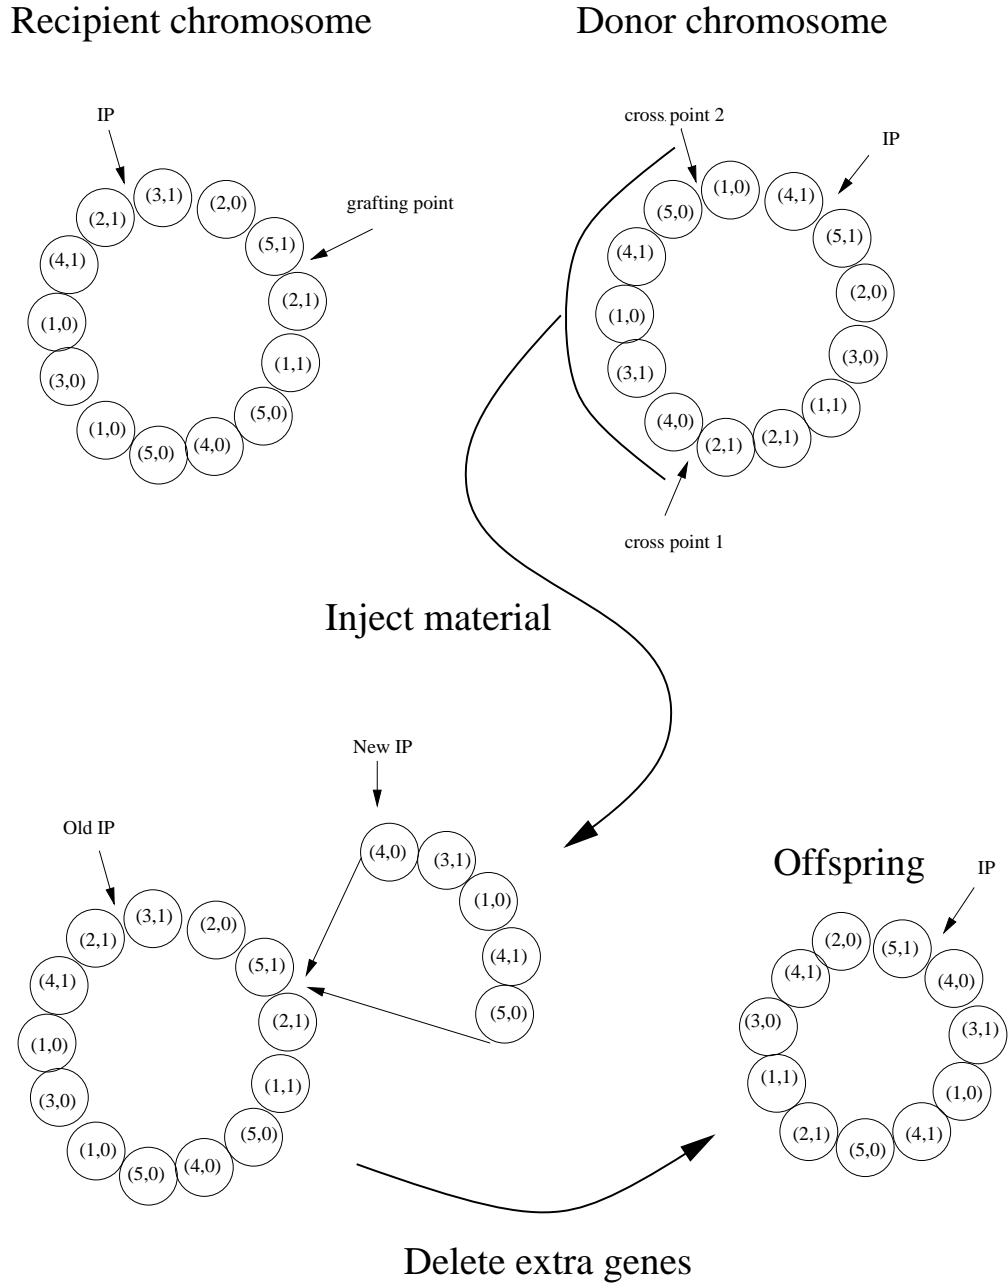


Figure 2: The crossover operator of the LLGA. Step 1: pick a grafting point in the recipient, and pick a random segment from the donor. Step 2: inject the random segment at the grafting point of the recipient. Step 3: traverse the recipient and delete extra genes in order to make a legal EPE-2 chromosome.

## 4 Simulating the LLGA with compressed introns

To cope with the exponential growth of non-coding material, we introduce a technique that compresses the non-coding genes, thus making it tractable from a computational perspective. The idea is simple. Instead of keeping the non-coding genes individually, we just keep the size of the non-coding segment (the number of consecutive non-coding genes) between each pair of coding genes. The storage requirement becomes  $O(\text{number of coding genes})$ . The example in figure 3 illustrates the compressed representation.

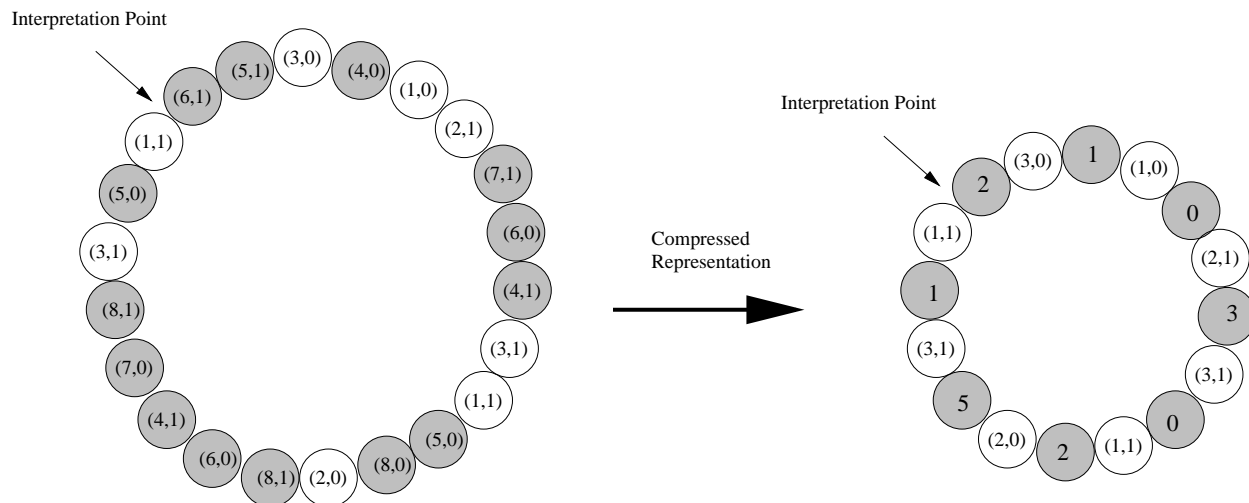


Figure 3: Discrete (left) and compressed (right) representation of non-coding genes. In this example, genes 1,2,3 are coding genes. Genes 4,5,6,7,8 are non-coding and are shown in grey for easy visualization.

With the compressed representation, the identity of the non-coding genes is lost. But this shouldn't be a problem because these genes give no contribution to the fitness function. Their purpose in the chromosome is to provide spacings to reduce the chance of BB disruption and to help the formation of linkage.

Having the compressed representation is just a trick to be able to cope with problems that require a large number of non-coding genes. But we also need to be able to simulate the LLGA's crossover operator as if it had the discrete coding. The expression mechanism remains the same, but the deletion of the extra non-coding genes needs to be simulated. An empirical analyses of the LLGA shows that very few genes are deleted from the grafted material. This suggests the implementation of the following deletion scheme:

1. Let  $n$  be the number of non-coding genes in the grafted material.
2. Don't delete any genes from the grafted material.
3. Delete  $n$  non-coding genes uniformly from the rest of the chromosome.

Experiments with the compressed representation and the new deletion procedure were performed. The modified version mimics the original LLGA. Figure 4 is a replication (both with discrete and compressed representation) of the 19 exponentially scaled BB problem that is reported in Harik's dissertation (pp 89-90). The graphs show the evolution of convergence and linkage of each BB. The linkage measure was defined by Harik (1997). It is a numerical value that indicates how tight a BB is. The tighter the BB, the closer the linkage is to the 1.

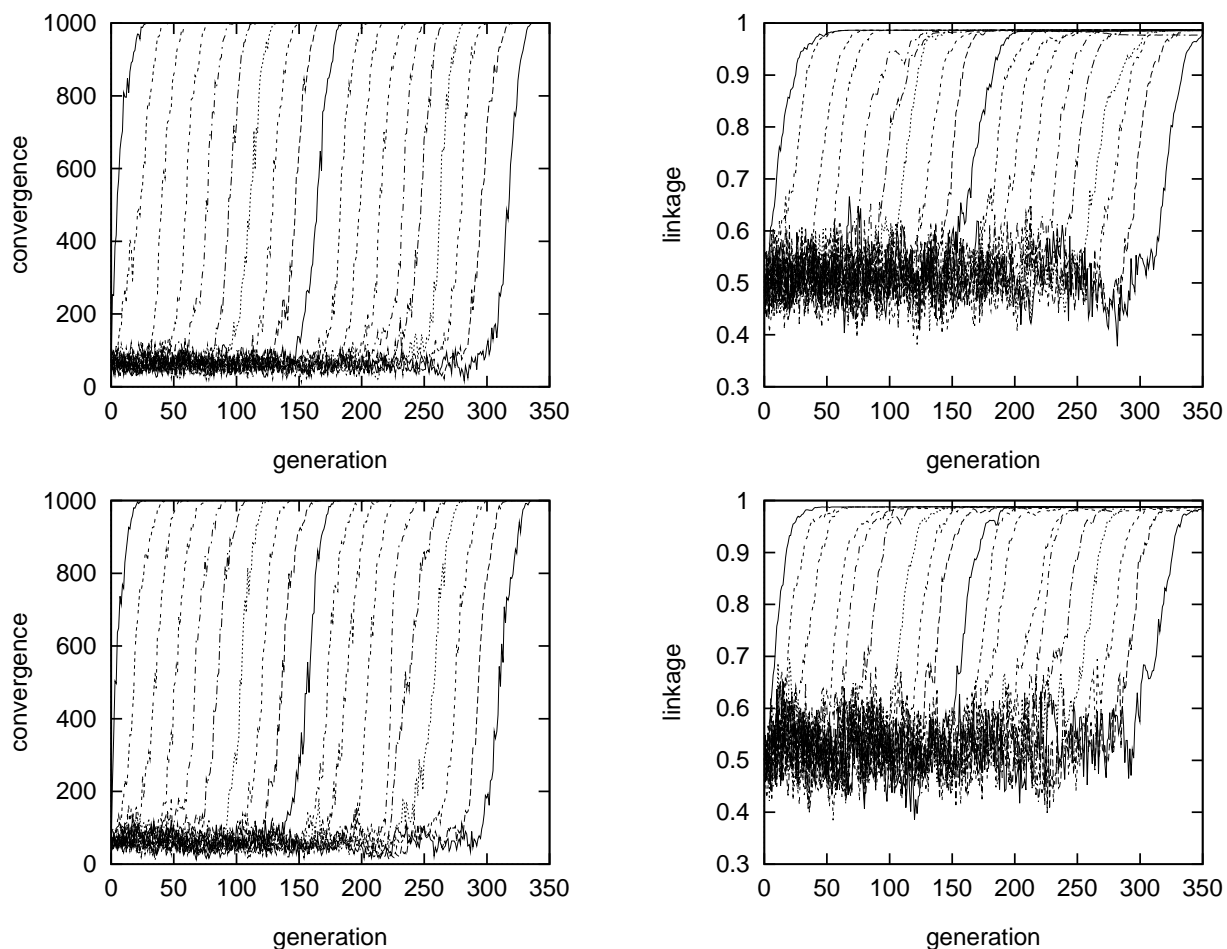


Figure 4: The LLGA with compressed and discrete representations of non-coding genes. The top graphs are for the discrete and the bottom ones for the compressed representation. The left graphs plot the convergence of each BB versus generation. The right graphs plot the evolution of linkage for each BB. Each BB is a 4-bit trap function with deceptive-to-optimal ratio of 0.6. The population size is 1000, and tournament selection of size 4 is used.

## 5 Time complexity of the LLGA on exponentially scaled problems.

This section presents an empirical study of the time complexity of the LLGA on exponentially scaled problems. We conduct experiments on problems with 20, 50, and 100 building blocks. Each consists of a 4-bit trap function with a deceptive-to-optimal ratio of 0.8 (Deb & Goldberg, 1993). We record the minimum population size needed to solve the problem to optimality at least 4 out of 5 independent runs. For all these problems, the LLGA takes about the same time to solve each BB (around 17 generations each). For 100 BBs, it takes near 1700 generations to solve all of them. The pattern is the same as the one observed in figure 4. Figure 5 is a log-log plot of the number of function evaluations needed to solve the problem to optimality as a function of the number of building blocks. The curve has a slope of approximately 1.3, which indicates that the time complexity of the LLGA is almost linear for exponentially scaled problems. The number of generations needed to solve a larger problem increases linearly, but the population size remains



nearly constant. The LLGA solves one BB at a time while the others wait for their turn. Diversity is never lost due to the probabilistic expression mechanism. For this type of problems, the initial supply of BBs dominates the population sizing. Correct decision-making is secondary since there's almost no noise coming from the other partitions (Goldberg, Deb, & Clark, 1992), (Harik, Cantú-Paz, Goldberg, & Miller, 1997). The GA can correctly decides between two individuals based on the outcome of the competition in one BB alone.

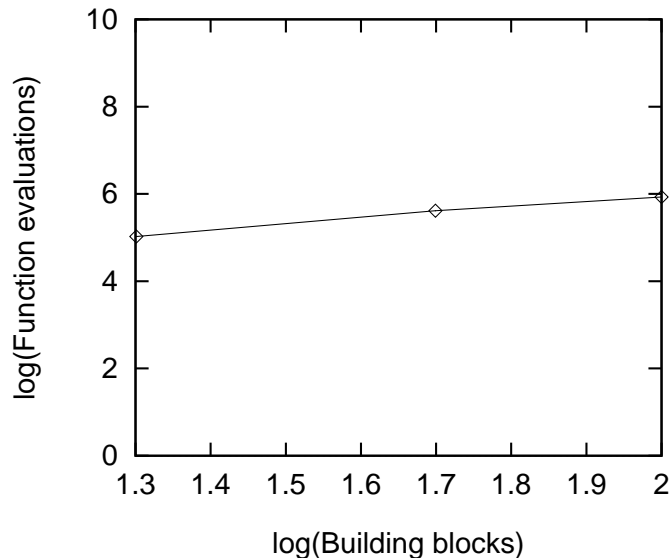


Figure 5: Log-Log plot of the number of function evaluations needed to solve the problem to optimality as a function of the number of building blocks. The slope of the curve is approximately 1.3, which means almost linear performance of the LLGA on exponentially scaled problems.

## 6 Conclusions

This paper revisited the linkage-learning genetic algorithm and showed that its time complexity is almost linear for exponentially scaled problems. Aside from that, this work represents a small step towards a competent GA. The efficient representation of the non-coding material is likely to be crucial for the solution of uniformly scaled problems within the linkage learning GA framework. Investigations in this direction have already begun and will be reported soon.

For the past 20 years, many applications have had success using SGA-like technology. But it is possible to do better. Competent GAs are coming soon, and with them, hard optimization problems will be able to enjoy success as well.

## Acknowledgments

This study was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grants F49620-94-1-0103, F49620-95-1-0338, and F49620-97-1-0050. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U. S. Government.

## References

- Deb, K., & Goldberg, D. E. (1993). Analyzing deception in trap functions. In Whitley, L. D. (Ed.), *Foundations of Genetic Algorithms 2* (pp. 93–108). San Mateo, CA: Morgan Kaufmann.
- Goldberg, D. E., Deb, K., & Clark, J. H. (1992). Genetic algorithms, noise, and the sizing of populations. *Complex Systems*, 6, 333–362.
- Goldberg, D. E., Deb, K., Kargupta, H., & Harik, G. (1993). Rapid, accurate optimization of difficult problems using fast messy genetic algorithms. In Forrest, S. (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms* (pp. 56–64). San Mateo, CA: Morgan Kaufmann.
- Goldberg, D. E., Korb, B., & Deb, K. (1989). Messy genetic algorithms: Motivation, analysis, and first results. *Complex Systems*, 3(5), 493–530. (Also TCGA Report 89003).
- Harik, G. R. (1997). *Learning gene linkage to efficiently solve problems of bounded difficulty using genetic algorithms*. Doctoral dissertation, University of Michigan, Ann Arbor. Also IlliGAL Report No. 97005.
- Harik, G. R., Cantú-Paz, E., Goldberg, D. E., & Miller, B. L. (1997). The gambler’s ruin problem, genetic algorithms, and the sizing of populations. In Bäck, T. (Ed.), *Proceedings of the Fourth International Conference on Evolutionary Computation* (pp. 7–12). New York: IEEE Press.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Kargupta, H. (1996). The gene expression messy genetic algorithm. In *Proceedings of the Third International Conference on Evolutionary Computation* (pp. 814–819). New York: IEEE Press.
- Levenick, J. R. (1991). Inserting introns improves genetic algorithm success rate: Taking a cue from biology. In Belew, R. K., & Booker, L. B. (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms* (pp. 123–127). San Mateo, CA: Morgan Kaufmann.
- Levenick, J. R. (1995). Metabits: Generic endogenous crossover control. In Eschelman, L. (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 88–95). San Francisco, CA: Morgan Kaufmann.
- Thierens, D., & Goldberg, D. E. (1993). Mixing in genetic algorithms. In Forrest, S. (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms* (pp. 38–45). San Mateo, CA: Morgan Kaufmann.
- Wang, L. (1997). Personal communication.
- Wu, A., & Lindsay, R. K. (1995). Empirical studies of the genetic algorithm with non-coding segments. *Evolutionary Computation*, 3:2, 121–147.