

**A Simple Implementation of
the Bayesian Optimization Algorithm (BOA)
in C++ (version 1.0)**

Martin Pelikan

IlliGAL Report No. 99011
March 1999

Illinois Genetic Algorithms Laboratory
University of Illinois at Urbana-Champaign
117 Transportation Building
104 S. Mathews Avenue Urbana, IL 61801
Office: (217) 333-0897
Fax: (217) 244-5705

A Simple Implementation of the Bayesian Optimization Algorithm (BOA) in C++ (version 1.0)

Martin Pelikan
Illinois Genetic Algorithms Laboratory
104 S. Mathews Avenue, Urbana, IL 61801
University of Illinois at Urbana-Champaign
Phone/FAX: (217) 333-2346, (217) 244-5705
`pelikan@illigal.ge.uiuc.edu`

Abstract

The paper explains how to download, compile, and use the simple implementation of the Bayesian optimization algorithm (BOA) (Pelikan, Goldberg, & Cantú-Paz, 1998; Pelikan, Goldberg, & Cantú-Paz, 1999), version 1.0, written in C++. It provides the instructions for creating input files for the BOA to solve various problems with various parameter settings and for adding new test functions into the existing code. Outputs of an example experiment are discussed.

1 Introduction

The purpose of this paper is to give basic instructions for downloading, compiling, and using the version 1.0 of the implementation of the Bayesian optimization algorithm written in C++ that is publicly available at the IlliGAL anonymous ftp-site. The installation instructions are designed for UNIX operating systems. However, we suppose no major modifications are necessary and the source codes should be compiled under most operating systems with a number of different compilers without major problems.

We have tried to keep the implementation as simple as possible and yet sufficiently powerful to demonstrate the basic principle of the algorithm and either to reproduce the results of recent experiments or to produce new ones. Therefore, the code does not include many features discussed in the papers discussing the algorithm as the incorporation of prior information or other than a simple greedy search for constructing the network modeling the data. Although some low-level constructs are written in object-oriented C++, on a higher level we have avoided the use of object-oriented features so that the code is tractable even for the users grown on structural or functional languages.

The paper starts by providing the instructions for downloading and extracting the package including the source code and a few example input files. Section 3 explains how to compile the extracted source code. Section 5 discusses the features of the implementation or what has actually been implemented. In section 6, the format of input files and the description of parameters that it can specify are presented. Section 7 discusses the format of output files for a simple example experiment. A short description of the source files, the list of implemented test functions, and the instructions for plugging in new test functions can be found in Section 8.

2 How to Download and Extract the Source Code

The package with the source code and few example input files is available at the IlliGAL anonymous ftp site in `ftp://ftp-illigal.ge.uiuc.edu/pub/src/sBOA/C++/sBOA.tar.Z`.

After downloading the package (`sBOA.tar.Z`) into your directory, the files can be extracted by typing the following:

```
uncompress sBOA.tar.Z
tar xvf sBOA.tar
```

After extracting the files correctly, your directory should contain the following files containing the source code and a directory with the example input files with their outputs:

K2.cc	checkCycles.cc	graph.h	population.h	stack.cc
K2.h	checkCycles.h	header.cc	random.cc	stack.h
Makefile	computeCounts.cc	header.h	random.h	startUp.cc
args.cc	computeCounts.h	help.cc	recomputeGains.cc	startUp.h
args.h	examples	help.h	recomputeGains.h	statistics.cc
bayesian.cc	fitness.cc	main.cc	replace.cc	statistics.h
bayesian.h	fitness.h	memalloc.h	replace.h	utils.cc
binary.h	getFileArgs.cc	mymath.cc	sBOA.tar	utils.h
boa.cc	getFileArgs.h	mymath.h	select.cc	
boa.h	graph.cc	population.cc	select.h	

In the directory `examples`, the files starting with `input` are example input files and the files starting with `output` are the outputs to these files. The name of these files says what fitness function these files are intended to optimize and what is the size of this function. Parameters that can be specified in the input files as well as the format of output files are explained in the remainder of this paper.

3 How to Compile the Source Code

The compilation is very simple. In a file `Makefile` that can be found in the directory with the uncompressed source code, the following two changes should be performed:

Line 36

In the statement `CC = CC`, the `CC` on the right-hand side should be changed to the name of a preferred C++ compiler on your machine. With a `gcc`, for instance, the line should be changed to `CC = gcc`.

Line 42

In the statement `OPTIMIZE = -O3`, the required optimization level should be set (for GNU `gcc` this is `-O4`, for SGI `CC` it is `-O3`). For no code optimization, use only `OPTIMIZE = .` For instance, for a maximal optimization with `gcc`, i.e. `-O4`, use `OPTIMIZE = -O4`. All modules are compiled at once since some compilers (as SGI `CC`) use intermodule optimization that does not allow them to compile each source file separately and link them together afterwards.

After making the above modifications to the `Makefile`, you can compile the simple BOA by typing the following line:

```
make all
```

We have tested the code for various operating systems and compilers. The results are summarized in Table 1. For some operating systems and compilers, problems with the code optimization might be encountered (see Table 1). Since the optimized code works much faster, we encourage you to try both alternatives (with and without optimization) and compare the first few generations of the runs (e.g., average fitness values) for both executables with the same random seed and input parameters. We have not identified the reason why our implementation gives different results with optimized and non-optimized code on several system-compiler combinations. However, giving different results does not necessarily imply that the optimized code is incorrect; it only means that the optimized code acts slightly differently. Nevertheless, the overall results are approximately the same.

Table 1: Compilation of the package on various systems and with several different compilers

Operating System	Compiler	Problems Encountered
IBM AIX 3.2	gcc 2.7.2	Optimized code produces different results.
IBM AIX 4.1	xlC 3.1	none
IBM AIX 4.2	gcc 2.7.2	Optimized code produces different results.
Linux 2.0.35	gcc 2.7.2	Optimized code produces different results.
Linux 2.1.123	egcs 2.90.27	Optimized code produces different results.
SGI Irix 6.2	CC 7.1	none
Sun OS 4.1.3	gcc 2.7.2	none
Sun OS 5.6	gcc 2.7.1	none
ULTRIX 4.4	gcc 2.7.2	none

After correctly compiling the source code there should be an executable file `boa` in your directory.

4 Command Line Parameters

Without any command line parameters, the `boa` runs with all parameters set to their default values. The `boa` can be called with either of the following parameters (for examples see Table 2):

`<filename>`

Runs the BOA with input parameters specified in the file `<filename>`.

`-h`

Prints the description of command line parameters.

`-paramDescription`

Prints the description, the type, and the default value of all input file parameters.

Table 2: Examples on command line parameters

Command line	Description
<code>boa</code>	runs the boa with all input parameters set to their default values
<code>boa myInputFile</code>	runs the boa with input parameters specified in the file <code>myInputFile</code>
<code>boa -paramDescription</code>	prints the description of all input file parameters
<code>boa -h</code>	prints the help on the command line parameters

5 What Has Been Actually Implemented?

The following list briefly summarizes what can be found in the version 1.0 of the BOA implementation:

Representation of Solutions

Solutions are represented by binary strings of fixed length.

Test Functions

Few decomposable test functions with and without overlapped building blocks have been implemented. The user can add his own test functions easily. Each test fitness function can (but need not) include the initialization and done methods. This feature can be useful for more sophisticated functions that need to read input parameters or allocate some memory before their first evaluation and perform certain actions in order to clean the used memory or the like after their last evaluation in the run).

Problem Size

The problem size can be set by the user. Very big problem sizes are allowed (up to 32767). The bigger the problem, the bigger the population, and the longer the run takes.

Population Size

The population size can be specified by the user. Very big populations can be used (up to 2,147,483,647). However, the bigger the population size, the slower it takes to process one generation and the more memory the algorithm uses, and therefore very big populations are not very useful. We have used populations up to 50,000.

Selection method

Truncation selection (also called block selection) has been implemented. Truncation selection selects the best portion of the population. The user can control the selection pressure by choosing the number of parents to select (in percent of population).

Replacement method

Replacement of the worst has been implemented. With replacement of the worst, the worst solutions in the original population are replaced by offspring. The user can specify the number of offspring (in percent of population).

Scoring Metric

The K2 metric derived from the Bayesian Dirichlet metric has been implemented. The metric is used as a measure in order to construct the network that is a good model for the set of selected promising solutions. It gives preference to networks that model this set more

accurately. For the sake of keeping the implementation simple, no prior information in form of the prior network or the set of high-quality solutions can be incorporated into the metric in our implementation.

Prior Information

No prior information but the maximal number of incoming edges into each node can be used. This number corresponds to the maximal order of interactions to be considered in the distribution estimate.

Network Construction Method

A greedy algorithm with only edge addition allowed has been implemented. The algorithm gradually adds the edges which increase the scoring metric the most until no edges increase the metric or no more edges can be added without breaking the constraint on a maximal number of incoming edges.

Output Statistics

There are few different outputs from the algorithm. The evolution of the best, average, and worst fitness values, the best solution in a current generation, the bias of the population, and the model used to generate offspring during the run can be extracted. All outputs can be related to the current number of generation or a number of fitness evaluations performed so far.

Termination Criteria

The run can be terminated after a maximal number of generations, maximal number of fitness evaluations, or a maximal proportion of optima in a population are reached. The run can also be stopped when the population has almost converged and the bits on all positions are almost homogeneous or when the optimum has been found. Any of the criteria can be ignored and any combination of various criteria can be used. If a termination criterion that uses a proposition that decides whether the solution is optimal or not, if this proposition is not defined (when the algorithm does not know what is the optimum and what is not), the criterion is ignored.

6 Input files

Input files can contain the statements of the following form:

```
<identifier> = <value>
```

where `<identifier>` is an identifier of a particular parameter and `<value>` is its new value. Empty lines and extra spaces that are not within the identifier or its value are ignored. The order of statements in the input file is not important. Each parameter can be defined at most once. If the value of a parameter is not specified in the input file, its default value is substituted. In the case of multiple definition of any parameter, the program ends up with an error message informing what parameter was multiply defined. If the identifier does not exist, the program ends up with an error message. The interpreter of input file is case sensitive.

The following list describes the values of parameters that can be specified in the input file, their types, and their default values. You can get a similar list by running `boa -paramDescription`.

`populationSize`

Description: The size of a population.
Type: `long`
Default: `1200`

`parentsPercentage`

Description: The number of parents to select (in percent of population).
Type: `float`
Default: `50`

`offspringPercentage`

Description: The number of offspring to generate (in percent of population).
Type: `float`
Default: `50`

`fitnessFunction`

Description: The number of a fitness function to use. (See Section 8.2 for the list of test functions included in the implementation).
Type: `int`
Default: `2` (deceptive of order 3 without overlapping)

`problemSize`

Description: The size of a problem (string length).
Type: `int`
Default: `30`

`maxNumberOfGenerations`

Description: Maximal number of generations to perform before terminating the run. `-1` stands for unlimited.
Type: `long`
Default: `200`

`maxFitnessCalls`

Description: Maximal number of fitness calls before terminating the run. `-1` stands for unlimited.
Type: `long`
Default: `-1` (unlimited)

`epsilon`

Description: A threshold for univariate frequencies for terminating the algorithm due to the so-called bit-convergence. If frequencies of all bits are closer than epsilon to either 0 or 1, the run is terminated. `-1` stands for ignoring this criterion.
Type: `float`
Default: `0.01`

`stopWhenFoundOptimum`

Description: Terminate the run when the optimum has been found (if the proposition identifying the optimum for the optimized function is defined)? A non-zero value stands for “Yes”, zero stands for “No”.

Type: `char`

Default: 0 (no)

`maxOptimal`

Description: Terminate the run when the proportion of optimal solutions (in percent of a population) has reached this value. `-1` stands for ignoring this criterion.

Type: `float`

Default: `-1` (ignore)

`maxIncoming`

Description: Maximal number of incoming edges into any of the nodes in the considered networks (denoted by k in the algorithm description). Corresponds to the maximal order of interactions that can be covered by the used class of models (it is equal to the order of interactions that can be covered minus 1).

Type: `int`

Default: 2 (interactions of 3. order)

`pause`

Description: Wait for Enter key after each generation? A non-zero value stands for “Yes”, zero stands for “No”.

Type: `char`

Default: 0 (no)

`outputFile`

Description: The base of output file names (will add the extensions to each output file name according to the type of the file).

Type: `char*`

Default: `NULL` (no output file)

`guidanceTreshold`

Description: A threshold for a population bias displayed each generation. As soon as the frequency of a bit gets closer than this parameter to either 0 or 1, the bit is said to be biased to the corresponding value.

Type: `float`

Default: 0.3

`randSeed`

Description: A random seed.

Type: `long`

Default: `time` (current time)

An example of input file is presented in Figure 6. With this input file the deceptive function of order 3 of the size (number of bits) 30 will be optimized with the population size of 1200. Truncation

selection that selects the best half of the solutions will be used. Each generation, the half of the original population is replaced by the offspring. The networks are to have two incoming edges into each node at maximum. The run will be terminated after either 300 generations are performed or the frequencies of all bits are closer than 0.01 to either 0 or 1. Outputs will be written to output files with the base `output.3deceptive.30` and additional extensions corresponding to their type. Random seed is set to 123. Other parameters are set to their default values. The presented input file is included in the package along with the output files.

```
populationSize = 1000
problemSize    = 30
fitnessFunction = 2

parentsPercentage = 50
offspringPercentage = 50

maxIncoming = 2

maxNumberOfGenerations = 40
epsilon                = 0.01

outputFile = output.3deceptive.30

randSeed = 123
```

Figure 1: Example input file (included in the package as `input.3deceptive.30`).

7 Outputs

In this section, we will briefly describe the outputs of the `boa` for the input file `input.3deceptive.30` shown in Figure 6 included in the package. We divide the section into two parts. The first one discusses the outputs to the standard output (which is mostly the screen). The second one discusses the outputs that can be found in produced output files (if any).

7.1 What Can You See on the Screen?

At first, the header with the name of the program, the author, the date of its release, and the name of input file is printed on the screen. It is followed by the list of all parameters and their values (not only those specified in the input file). For each parameter, its description, identifier, type, and the current value is displayed. Since this initial output described above is very simple and easy to understand, we do not present an example here.

After printing the information about the product and the parameters, the information about the generation number, the number of fitness evaluations performed so far, the best, average, and the worst fitness values in a current population, the proportion of optima in a current population (if the proposition checking for optima is defined for the used test function), the population bias (the guidance vector), and the best solutions in the current population is displayed. This information is printed out each generation. An example of the output information written each generation is shown in Figure 7.1. The example was produced with the `boa` with input parameters specified in

the input file `input.3deceptive.30`, included in the package. It shows, that the number of current generation is 10, 5500 fitness evaluations were performed so far from the beginning of the run, the best fitness in the current population is 10, the average fitness in the population is 9.6337, and the worst fitness in the population is 8.900001. There are 1% of global optima in the population. The bits on positions from 1 to 12 and from 16 to 24 are biased to 1 (the frequencies of 1 on these positions is closer than the parameter `guidanceTreshold` to 1). All other positions are unbiased (the “.” is displayed on unbiased positions). The best solution in a population has 1’s on all positions.

```

Generation           : 10
Fitness evaluations   : 6000
Fitness (max/avg/min) : (10.000000 9.633700 8.900001)
Percentage of optima in pop. : 1.00
Population bias       : 111111111111...1111111111.....
Best solution in the pop. : 11111111111111111111111111111111

```

Figure 2: Example information produced by the `boa` each generation to standard output.

After the run is terminated, the similar information is produced for the last generation. In addition to this, the information on the reasons for terminating the run is provided (in case of meeting multiple criteria for terminating the run, the first identified reason is displayed). Again, when the proposition for checking the optima is not available, the items that use this are excluded. An example of the final information closing the run is shown in Figure 7.1. The example was produced with the `boa` with input parameters specified in the input file `input.3deceptive.30`, included in the package.

```

FINAL STATISTICS
Termination reason      : Bit convergence (with threshold epsilon)
Generations performed    : 15
Fitness evaluations      : 8500
Fitness (max/avg/min)    : (10.000000 10.000000 10.000000)
Percentage of optima in pop. : 100.00
Population bias          : 11111111111111111111111111111111
Best solution in the pop. : 11111111111111111111111111111111

```

Figure 3: Example information produced by the `boa` at the end of the run.

After performing the whole run, a string “The End.” should be printed out, as the last line sent to the standard output.

7.2 What Can You Find in the Output Files?

If you specify the parameter `outputFile` in the input file, the `boa` produces three output files, each starting with the base given by the `outputFile` parameter with an extension depending on the file type.

The following list displays the file names and the short description of the content corresponding output files with respect to the `<base>` (the base of the file names given by the `outputFile`

parameter):

<base>.log

Log-file, including all standard outputs (except for the message about waiting for the Enter key to be pressed).

<base>.model

The models (the description of the used network) used for generating offspring each generation.

<base>.fitness

The best, average, and the worst fitness, with the number of current generation and the number of fitness evaluations performed in a format that is easily visualized by visualization tools as **gnuplot**.

The log-file includes all standard outputs that are explained in the previous section and therefore no additional explanation is necessary.

In the file with the networks used to generate the offspring each generation, the number of current generation and the list of nodes and their parents are displayed each generation. Each node is denoted by the number of variable it corresponds to starting with 0. Each line providing the information about the node and its parents is of the following form:

```
<node> <- <list of its parents>
```

Parents are divided by commas. If there are no parents, the right-hand side remains empty. An example of the output in this file is shown in Figure 7.2. The example was produced with the **boa** with input parameters specified in the input file **input.3deceptive.30**, included in the package. In our example, in the network used to generate the offspring in generation 5 there are two edges ending in the node 0, the first starts in 1 and the second starts in 2 (i.e., the node 0 has two parents, nodes 1 and 2). Similarly, for instance, there is one edge ending in the node 9. This edge starts in the node 5. The node 5 has no parents at all, i.e. there are no edges ending in this node. Remaining information from the example can be interpreted in a similar way.

In the file with the extension **fitness** including the best, worst and average fitness values with respect to the number of current generation and the number of fitness evaluations performed so far, each line includes 5 numbers. The following list describes the values (in the order from the leftmost)

- 1) The number of current generation
- 2) The number of fitness calls performed so far
- 3) The best fitness in a current population
- 4) The average fitness in a current population
- 5) The worst fitness in a current population

An example of a line produced by the **boa** with a specified output file name with additional extension **fitness** is shown in Figure 7.2. The example was produced with the **boa** with input parameters specified in the input file **input.3deceptive.30**, included in the package. The line says that in generation 7, after performing 4500 fitness calls, the best fitness in the population is 9.8, the average fitness is 9.1862, and the worst fitness is 8.400001 (this is an error caused by floating-point operations, the correct value is 8.4).

```
Generation: 5
```

```
0 <- 1, 2
1 <- 2, 10
2 <- 4, 11
3 <- 5, 10
4 <- 3, 5
5 <-
6 <- 8, 12
7 <- 6, 8
8 <- 9, 26
9 <- 5
10 <- 11, 9
11 <- 9, 5
12 <- 15, 29
13 <- 12, 14
14 <- 12, 18
15 <- 16, 17
16 <- 17, 22
17 <- 11, 2
18 <- 20, 19
19 <- 9, 2
20 <- 19, 29
21 <- 19, 0
22 <- 21, 23
23 <- 21, 4
24 <- 25, 26
25 <- 10, 19
26 <- 25, 28
27 <- 0, 17
28 <- 27, 16
29 <- 27, 28
```

Figure 4: Example model description produced by the `boa`.

8 The Code

In this section, we shortly describe the function of each of the source files. Thereafter, we provide the instructions for plugging in a new test function into the existing code.

8.1 Brief Description of the Source Files

The following list briefly describes what functions are located in each source file. A similar description is located at the beginning of the corresponding source files among with the information about the author and the date of a last modification. For each function, a detailed description of its purpose and its input parameters are presented before its definition. The source files are heavily commented.

K2.cc

Functions for the initialization and use of the K2 metric.

7	4500	9.800000	9.186200	8.400001
---	------	----------	----------	----------

Figure 5: Example line in the output file including the information about the fitness.

`args.cc`

Functions for manipulation with arguments passed to a program.

`bayesian.cc`

Functions for construction and use of Bayesian networks (not including the metric related functions and some of more specific functions defined elsewhere).

`boa.cc`

Functions for the initialization of the BOA, the BOA itself and a done method for the BOA.

`checkCycles.cc`

A function that checks what edges would create cycles with a newly added edge and assigns negative gain for the corresponding edge additions.

`computeCounts.cc`

Functions for computing counts for all instances of a particular list of string positions or a list of positions created by adding a position from the array of positions to a particular list of remaining positions repeatedly.

`fitness.cc`

The definition of fitness functions; in order to add a fitness one has to add it here (plus the definition in the header file `fitness.h`); see documentation or the instructions below.

`getFileArgs.cc`

Functions for reading the input file, printing the description of the parameters that can be processed, and the related.

`graph.cc`

The definition of classes `OrientedGraph` and `AcyclicOrientedGraph` for manipulation with oriented graphs.

`header.cc`

Prints out the header saying the name of the product, its author, the date of its release, and the file with input parameters (if any).

`help.cc`

help (arguments description, input file parameters description).

`main.cc`

Main routine and the definition of input parameters.

`mymath.cc`

Commonly used mathematical functions.

`population.cc`

Functions for manipulation with the populations of strings and the strings themselves.

`random.cc`

Random number generator related functions (random generator is based on the code by Fernando Lobo, Prime Modulus Multiplicative Linear Congruential Generator (PMMLCG)).

`recomputeGains.cc`

A function calling the metric repeatedly in order to recompute the gains for all edge additions ending in a particular node.

`replace.cc`

The definition of replacement replacing the worst portion of the original population and the divide and conquer function it uses to separate the worst.

`select.cc`

The definition of truncation selection and the divide and conquer function it uses to separate the best.

`stack.cc`

The definition of a class `IntStack` (a stack for `int`).

`startUp.cc`

A start-up function for processing the arguments passed to the program and the function returning the name of the input file if any was used.

`statistics.cc`

Functions that compute and print out the statistics during and after the run.

`utils.cc`

Functions use elsewhere for swapping values of the variables of various data types.

8.2 Implemented Test Functions

We have implemented the following test functions (ordered by their number).

Number	Description
0	OneMax (bit-count) function
1	Quadratic function without overlapping
2	A deceptive function of order 3 without overlapping
3	A trap function of order 5 without overlapping
4	A bipolar function of order 6 without overlapping
5	A deceptive function of order 3 with 1-bit overlap between adjacent building blocks

For the definition of the quadratic function and the trap function of order 5, see Pelikan and Mühlenbein (1999). For the definitions of the rest of the functions, see Pelikan et al. (1998) or Pelikan et al. (1999). In the latter two papers, there is a typo in the definition of the trap function.

8.3 How to Plug-in a New Test Function

The instructions to plug in a new test function, also provided in `fitness.cc` follow:

1. Create a function with the same input parameters as other fitness functions defined in this file (e.g., `onemax`) that returns the value of the fitness given a binary chromosome of a particular length (sent as input parameters to the fitness). Place the function in the source file `fitness.cc`
2. Put the function definition in the `fitness.h` header file (look at `onemax` as an example).
3. In file `fitness.cc`, increase the counter `numFitness` and add a structure to the array of the fitness descriptions `fitnessDesc` as described below. For compatibility of recent input files, put it at the end of this array in order not to change the numbers of already defined and used functions. The structure has the following items (in this order), we also provide an example in the form of how the items are set for `onemax` function:
 - a) a string description of the function (informative in output data files). For `onemax` the description is "ONEMAX".
 - b) a pointer to the function (simple "&" followed by the name of a function). For `onemax` this is `&onemax`, since this function is defined in a function named `onemax`.
 - c) a pointer to the function that returns true if an input solution is globally optimal and false if this is not the case. If such function is not available, just use `NULL` instead. The optimum of `onemax` is in 111...1 and therefore the function `areAllGenesOne` (which returns true if the input string has 1's on all positions) is used. This item is therefore set to `&areAllGenesOne` with `onemax` function.
 - d) a pointer to the function for initialization of the particular fitness function (not used in any of these and probably not necessary for most of the functions, but in case reading input file would be necessary or so, it might be used in this way). Use `NULL` if there is no such function. In `onemax`, no initialization is necessary and therefore this item is set to `NULL`.
 - e) a pointer to the "done" function, called when the fitness is not to be used anymore, in case some memory is allocated in its initialization; here it can be freed. Use `NULL` if there is no need for such function. In `onemax`, no such function is necessary and therefore this item is also set to `NULL`.
4. The function will be assigned a number equal to its ordering number in the array of function descriptions `fitnessDesc` minus 1 (the functions are assigned numbers consequently starting at 0); so its number will be equal to the number of fitness definitions minus 1 at the time it is added. Its description in output files will be the same as the description string (see 3a).

9 Final Comments

The code is distributed for academic purposes with absolutely no warranty of any kind, either expressed or implied, to the extent permitted by applicable state law. We are not responsible for any damage resulting from its proper or improper use.

If you have any comments or identify any bugs, please contact the author (email is a preferred way of communication).

Acknowledgments

The author would like to thank David E. Goldberg, Erick Cantú-Paz, Fernando Lobo, and Pavel Petrovic for their valuable comments to the paper as well as the implementation itself.

The work was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant number F49620-97-1-0050. Research funding for this project was also provided by a grant from the U.S. Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement DAAL01-96-2-0003. The author was also supported by grants number 1/4209/97 and 1/5229/98 of the Scientific Grant Agency of Slovak Republic.

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 mine and should not be interpreted as necessarily representing the official policies and endorsements, either expressed or implied, of the Air Force of Scientific Research or the U.S. Government.

References

- Pelikan, M., Goldberg, D. E., & Cantú-Paz, E. (1998). *Linkage problem, distribution estimation, and Bayesian networks* (Technical Report 98013). Urbana, IL: University of Illinois at Urbana-Champaign.
- Pelikan, M., Goldberg, D. E., & Cantú-Paz, E. (1999). *Boa: The bayesian optimization algorithm* (Technical Report 99003). Urbana, IL: University of Illinois at Urbana-Champaign.
- Pelikan, M., & Mühlenbein, H. (1999). The bivariate marginal distribution algorithm. In Roy, R., Furuhashi, T., & Chawdhry, P. K. (Eds.), *Advances in Soft Computing - Engineering Design and Manufacturing* (pp. 521–535). London: Springer-Verlag.