

# Applying Genetic Algorithms to Land Use Planning.

K.B. Matthews<sup>1</sup>, Susan Craw<sup>2</sup>, I. MacKenzie<sup>2</sup>, S. Elder<sup>3</sup> and A.R. Sibbald<sup>1</sup>

**Abstract.** This paper explores the potential of applying Genetic Algorithms to land use planning, a spatial allocation problem. Two genotype representations are proposed: a fixed-length genotype composed of genes that map directly to a land parcel's use, and a variable-length, order-dependent representation making allocations indirectly via a greedy algorithm. The fixed-length genotype is used within a standard genetic algorithm framework but the variable-length genotype requires novel breeding operators to be defined and post-processing of the genotype structure to identify and remove duplicate genotypes. The two approaches are compared on a real land use planning problem and the strengths and weaknesses of each approach are identified.

**Key Words:** representation, messy GAs, non-fitness information, land use planning

## 1 Introduction

Land use planning is a spatial allocation problem, where the planner, by manipulating the proportions and locations of land uses, seeks to satisfy one or more goals. Land use planning is a potentially challenging search and optimisation task, as the planner must frequently take into account complex non-linear interactions between parcels of land allocated to particular land uses.

GAs have several potential advantages as land use planning tools over more traditional methods. While GAs cannot be guaranteed to find an optimum solution [6], their efficiency in finding "good enough" solutions is well known [8]. In the case of a land use plan, allocations resulting in gains over the current land use define a minimum level of utility, while ability to generate allocations equivalent to those of an expert advisor is the goal. When used practically, it is also likely that the "optimum" plan found by the GA will form the strategic basis of an implemented plan, with further practitioner customisation to meet tactical management needs, rather than being implemented as it stands. The goal in most land use planning applications is to support rather than make decisions.

The population basis of GA search is particularly useful in land use planning since, depending on the replacement strategy adopted, it is possible to find a range of good candidate solutions. Examination of the differences between these solutions can lead to an enhanced understanding of key locations and processes.

GAs' utility as the foundation of flexible tools, capable in a range of application domains, is enhanced

by the independence of the GAs' search and optimisation mechanisms from their evaluation methods. GAs' robust ability to find acceptable solutions in complex search spaces also adds credibility to the solutions found. Exploiting the flexibility and robustness of GA search is, however, dependent on the existence of an appropriately responsive model to evaluate candidate solutions, the effective design of a representation for the elements of the application and a compatible well-parameterised set of operators [4].

This paper presents a comparison of two land-use planning GAs with common underlying structure but contrasting representation and operator sets. Section 2 examines the origins of the elements integrated within this research. Section 3 gives an overview of the land use planning domain and the wider research project of which the GA application is a key part. Section 4 details the GA representations considered and the two contrasting representations being pursued. Section 5 presents the operator set and the use of non-fitness information. Section 7 examines the experimental results achieved and the visualisations used to explore the features of the terminal population. Section 8 summarises the findings.

## 2 Related Work

The common underlying GA structure is based on Davis' Object-Oriented GA [4], employing fixed size, unstructured populations, with genotype uniqueness enforced and individual replacement of the lowest-fitness genotypes by offspring. Selection is rank-based, utilising a linear fitness normalisation function [15].

Explicitly spatial representations [2] were considered, but these were rejected as impractical for the land use planning application because they increased the magnitude of the optimisation problem and introduced further significant problems in translating the optimum solutions found into practical land use plans. Eshelman's analysis of building block disruption [7], however, made it possible to consider the alternative of treating land use allocation as a multiple-parameter, design optimisation problem using a one-dimensional genotype [1]. This led to the implementation of the *land-block* representation (Section 4.2).

Concern over limitations on the ultimate size of land allocation problems that could be successfully represented in this way led to further consideration of alternative representations and the utility of

<sup>1</sup> Land Use Science Group, Macaulay Land Use Research Institute, Aberdeen, K.Matthews@mluri.sari.ac.uk

<sup>2</sup> School of Computer and Mathematical Sciences, The Robert Gordon University, Aberdeen, S.Craw@scms.rgu.ac.uk IM@scms.rgu.ac.uk

<sup>3</sup> School of Electronic and Electrical Engineering, The Robert Gordon University, Aberdeen, S.Elder@eee.rgu.ac.uk

hybridising the GA with other algorithms. The best known examples of GA hybridisation are those with local search operators, to ensure the optimality of solutions once the general nature of the optimum solution has been found by the GA [8]. Cox et al., however, demonstrated a more general hybridisation approach, where a series of heuristics and algorithms were combined with the GA to achieve a solution [3]. An existing land allocation algorithm [10], whose operation is largely unaffected by the number of land-blocks to be allocated, was thus considered as the basis for a hybrid land allocation GA. This resulted in a hybrid GA employing a *percentage and priority* representation (Section 4.3).

Representing the candidate solutions for the percentage and priority GA, however, required a more flexible representation than used in the land-block GA. The information-rich, messy GA representation developed by Goldberg et al. [9], with its use of variable-length genotypes, order dependent interpretation of gene values, mechanisms for the elimination of redundant genes and ability to cope with both over and under-specified genotypes, provided the basis of the percentage and priority representation. The selected representation ultimately required GA operators capable of manipulating real-coded genes [11] and reordering of order-based genotypes [9].

### 3 Land Allocation Application

The GA-based land use planning tools are a key component within the Land Allocation Decision Support System (LADSS) being developed to assist land managers in making complex land management decisions.




Figure 1. LADSS Components

LADSS has four components in addition to the *GA-based land use planning tools*, (Figure 1). The *geographic information system* (GIS) provides the spatially-defined data (soils, climate and topography) and spatial analysis functionality required by the land use modules and the impact assessments. It also

provides useful facilities for the visualisation of the GA population as a series of maps.

The *land use modules* make individual, field scale, assessments of the suitability, productivity and profitability for a range of land uses. These estimates are based on biophysical data from the GIS, management parameters (e.g. the amount of chemicals applied to a crop) and global parameters (e.g. the market price for produce). It is worth noting that the suitability estimates are used by both the land-block gene initialisation and mutation methods (Section 5.2) to restrict the land allocations to those assessed as suitable. The productivity assessments are also used as the basis for sorting land-blocks within the percentage and priority GA allocation process (Section 4.3).

Information from individual fields is integrated for the land management unit as a whole by the *impact assessments*. The impact assessments are thus the fitness functions for the GA.

The *graphic user interface* displays maps of the land uses allocated by the GA during the course of a run in addition to the more usual GA visualisations. The use of the mapping facilities also allows the inclusion of a wealth of contextual information to aid the interpretation of the genotypes.

## 4 Genotype Representation

The land allocation application requires a genotype to encode the information needed to assign land uses to all land areas of interest to the planner. Three alternative representations were examined.

### 4.1 Spatially explicit representation

The first genotype representation considered was a two-dimensional grid of genes each determining the land use for a land parcel. The representation is frequently employed in spatial modelling [12], and was previously proposed, with two-dimensional operators by Cartwright and Harris [2]. As spatially adjacent land parcels are likely to have similar biophysical conditions, and thus correlated suitability and productivity values, it is likely that adjacent land parcels will form building-blocks from which an overall solution could be constructed. The second dimension of the grid-based representation, by mirroring in the genotype the adjacency of land parcels, minimises the disruption of the building-blocks by one- and two-point crossover operators.

The desire to use the more aggressively exploratory uniform crossover operator [14], enabled by adopting an individual replacement strategy for the GA, has the consequence of making building-block disruption position-independent [7]. This made it possible to adopt a conventional one-dimensional genotype representation with each gene representing a block of land, to which a land use would be allocated.

### 4.2 Land-block representation

This land-block representation has strong similarities

with that successfully adopted for other multiple-parameter optimisation problems [1]. The land-block representation is a fixed-length, fixed-order genotype made up of individual genes defining a parameter, in this case the land use for a single land-block (e.g. a field or forest compartment). Each gene is mapped to a land-block within the DSS, with the geometry of the land-block handled by the GIS as part of the evaluation process.

The land-block representation is superior to the spatially-explicit representations as these increase the magnitude of the optimisation problem. For example a farm with 100 fields can be represented by a land-block genotype with 100 genes. The same 100 genes would permit only a 10 by 10 grid over the area of the enterprise. This is clearly an inadequate spatial resolution for making land management planning decisions. While a quad-tree [13] representation would mitigate the increase in the number of elements required to represent a given land use allocation, at an acceptable spatial resolution, the reduction in data structure size is bought at the expense of a significant increase in its complexity.

The case for adopting the land-block representation is further strengthened when one considers that the land-blocks reflect real world enterprise infrastructures, with significant capital value, and have proven practical utility as management units, based on their size and biophysical characteristics. Further, while grid-based or quad-tree allocations can represent an optimum solution, the translation of these representations into land management units is an undesirable complication. Indeed, such a translation process could represent a significant optimisation problem in its own right.

A potential limitation of the land-block representation is its ability to cope with very large allocation problems. As genotype length depends directly on the number of land-blocks to be allocated, and problems have previously been encountered in GAs using very long genotypes [4], there can reasonably be expected to be a point at which the land-block GA will fail to find acceptable solutions. The typical size of allocation problem to date has 100 blocks and these have been successfully tackled. It was, however, decided to develop a second representation that would be largely independent of the number of land-blocks being considered.

### 4.3 Percentage and priority representation

This representation encodes the percentages of land to be allocated to each land use and the priority, or order, of their allocation. This is referred to as the percentage and priority (P&P) representation and makes genotype size dependent on the much smaller number of prioritised land-use percentages to be allocated, typically less than ten compared with 100 land-blocks.

There are three components to the representation: the land use, the percentage and the priority. The object class of the gene (represented as gene colour in Figure

2) determines the land use, the percentage is set by a real-coded parameter of the gene and the priority determined by reading the genotype in a fixed order (in the example from top to bottom). The ordered reading of the genotype has some similarities with the ICIS approach of Goldberg's messy GAs [9]. In Figure 2 the first P&P genotype exemplifies incomplete allocation (the target area percentages summing to less than 100%) and individual land uses appearing more than once. The genotype and gene structure is thus significantly more complex than for the land-block genotype with variable genotype length and gene evaluation depending on gene order. This is reflected in a substantially larger numbers of operators, Section 5.




Figure 2. P&P Representation

The allocation of land uses to specific land-blocks is handled by a "greedy algorithm" inspired by Cox [3]. The genotype is first decoded into a prioritised list of target land use percentages. The greedy algorithm then proceeds to iteratively allocate land-blocks, starting with those having the best performance per unit area for the highest priority land use, until the required percentage is exceeded or no land block for which the land use is suitable remain. The algorithm then proceeds to the next highest priority land use and continues until either all the land use percentages are achieved or no land-blocks remain to be allocated.

While the greedy algorithm and the P&P allocation combine to reduce the size of any land-allocation problem to the number of P&P genes required to specify an allocation (typically between 5 and 10), it does introduce a number of known biases into the allocations found. The greedy algorithm tends to over-allocate the higher priority land uses, as irregularly sized land-blocks do not combine to match target percentages. Choosing the land-block that resulted in the best fit with the target percentage, if a number of land-blocks with the same productivity per hectare were available for allocation, could reduce this over-allocation. The ordering of blocks for allocation is, however, the primary source of bias in the allocations. The use of productivity per hectare (Section 3), while logical and outwardly neutral as a metric for deciding the order in which blocks are allocated, could be

argued to bias the allocations in favour of production maximisation. This is acceptable for optimisations based on economic goals but may be an unacceptable bias in an environmental optimisation, and will skew results if economic and environmental factors were considered in the same multiple objective analysis. For multiple objective assessments it may be appropriate to include the ordering metric as the initial gene within the P&P genotype. Ultimately the choice of ordering metric(s) will be determined explicitly as part of the parameterisation of the land- use planning tools.

Random population initialisation and subsequent changes to gene priorities and percentages means that there can be no guarantee that the allocations defined in the genotype can be made, for example insufficient suitable land-blocks may remain. Since the failure to achieve the target percentage is detected by the DSS, it was decided to feed back to the GA information on the actual allocations made, in addition to the fitness evaluation. This allows the post-processing of P&P genotypes to deal with the following phenomena.

**Parasitic genes** occur when the value of the percentage attribute of a gene is revised to zero by the DSS (for example the second Wheat gene in Figure 2). This revision can occur as there are either no suitable, or simply no, land-blocks remaining to be allocated (an example of over specification). These genes are termed parasitic as they can make no contribution to the fitness of the genotype and survive only by virtue of the fitness of the other genes in the genotype. Genotypes containing these parasitic genes are functionally identical to a genotype without the parasitic genes as they result in the same allocation and evaluation by the DSS. The presence of these functionally identical genotypes is highly undesirable, since it reduces the effective genetic diversity of the population with consequent dangers of premature convergence. Zero-valued genes are thus removed after evaluation by the DSS and before testing for uniqueness and insertion into the population.

**Functionally identical genotypes** may also be created when the percentage values held in the genes are revised downward, by a failure to achieve the target percentage, or upward when the land-block allocation exceeds the target value. When this revision results in a genotype that already exists in the population it must be eliminated. This requires a second uniqueness test, applied after evaluation by the DSS.

**Gene pairs** occur when two genes of the same class, thus defining the same land use to be allocated, occur in consecutive gene loci (for example the two Forestry genes in Figure 2). This genotype is functionally identical to a genotype where the gene pair is replaced by a single gene whose percentage is the sum of the percentages of the gene pair. The gene pairs phenomenon is again undesirable for the same reasons as parasitic genes. Gene pairs are combined into single genes after the elimination of parasitic genes

**Incomplete allocation of land-blocks** (or under

specification) is dealt with by competition between genotypes, as a complete allocation will usually result in a fitter genotype. There is no need to complete the allocation using a competitive template [9] as it is possible to evaluate incomplete allocations.

## 5 Operators

We must now define the operators that provide appropriate exploration of the search space defined by the land-block and P&P representation. Table 1 shows the operator sets for the two representations.

**Table 1.** Operator sets for each representation

| Type   | Land-block        | Percentage and Priority                                                          |
|--------|-------------------|----------------------------------------------------------------------------------|
| Binary | Uniform Crossover | Uniform Crossover<br>Order Prioritised Crossover<br>Splice                       |
| Unary  | Mutation          | Type Mutation<br>Non-Uniform Mutation<br>Pair-Swap<br>Insert-Gene<br>Delete-Gene |

### 5.1 Binary Operators

Both representations employ uniform crossover [14], implemented using a crossover mask, with the crossover proportion set to maximise exploratory power [7]. This operator was chosen based on its record as a robust, high performance operator over a range of applications [4]. Since the P&P representation has variable genotype length the uniform crossover is performed only between the initial segment of the genotypes common to both parents.

The P&P representation also uses an order-prioritised crossover, (OPXO), based on uniform order-based crossover [4]. Uniform order-based crossover is used where permutations of gene order are significant. The operator first selects a subset of genes using a crossover mask. Subsequently the gene subset in each parent is reordered to the order of the equivalent genes in the other parent. The reordering is relative rather than absolute. While permutations of component genes are important to the P&P GA, only a subset of genes will be common to the parent genotypes. OPXO reorders the common genes to their order in the other parent genotype.

The final binary operator is splice, which concatenates two genotypes to form a double length genotype [9]. This provides a simple mechanism for increasing the number of elements in a candidate solution. The unary operator cut usually complements splice, but is not used with the P&P representation as it would result primarily in partial allocations, and these would be eliminated immediately by their inability to compete with existing completely allocated members of the population. Reductions in genotype length do, however, result from the elimination of parasitic genes and gene pairs. The number of genes within genotypes may also be modified by the unary mutation style operators gene-insert and gene-delete.

## 5.2 Unary Operators

For the land-block representation a single mutation operator is required. This replaces the current value of the land use of a gene with one chosen at random from the remaining suitable land use options.

The more complex P&P representation requires operators to mutate the land use, its target percentage and its priority. The land use is mutated by changing the class of the gene while preserving its position on the genotype and its value. This is termed type-mutation.

The target percentage is mutated using non-uniform mutation [11] of the real-coded value of the gene. The range of possible mutation values is constrained by upper and lower bounds and a granularity parameter carried by the individual genes. The granularity parameter allows the accuracy of specification for candidate solutions to be matched to the ability of the DSS to discriminate between them in its evaluation. Initial testing has used a granularity of 5%, though finer granularities will require experimentation.

The mutation of priorities is achieved by swapping the location of a pair of genes on the genotype. This pair-swap is based on the 2-opt operator described by Cox [3].

Two additional mutation operators change genotypes: one inserts a randomly generated gene into the genotype and the other deletes a single randomly selected gene. The insert-gene operator provides a mechanism to increase the number of individual land uses present in the land use plan. The delete-gene operator allows the promotion of lower-priority land use percentages that may be being “blocked” from achieving high fitness solutions by higher priority land uses. Adaption

## 6 Parameterisation

In both GAs, the operators were deployed using Davis’ independent operator-based reproduction with each offspring genotype being the product of a single operator [4]. The probability of applying an operator is adapted during the course of the run. Currently the balance of operator probabilities between the binary and unary operators is initially 0.65 to 0.35, adapting to 0.5 to 0.5 over the course of the run. Where more than one operator of a type exists, then the probabilities are equally divided among these operators. The adapting of operator probabilities responds both to how far towards the maximum length of the GA run the GA has progressed, and the number of reproductive events that have failed to result in a genotype making a fitness gain. This double-adaption results in a saw tooth pattern of adaption ensuring that even if the population converges well before the maximum run length the mutation operators are applied at the higher rate of probabilities before the run terminates.

The optimum balance of operator probabilities will, however, require to be established by a more thorough investigation into the parameterisation of both GAs.

## 7 Experimental Results

The goal of the initial testing was to establish the relative performance of the two representations on a typical land allocation problem.

### 7.1 Test Problem

The land allocation problem initially used to test the GAs was a research station in a disadvantaged area of the Scottish uplands. The research station is composed of 66 land-blocks, and 10 possible land uses are available. The evaluation used was the potential economic productivity of the research station (£ sterling). It should be noted that this represents a relatively simple optimisation problem, as it includes no impact assessment capable of accounting for the spatial interactions between land-blocks. The problem was useful as a test, however, as the global optimum may be simply determined.

### 7.2 Criteria Measured

To compare the quality of solutions found by the land-block and P&P GAs the fitness of the fittest member of the population (MaxFit) and the mean fitness of the population (AvgFit) was used.

To investigate the effectiveness of the learning by each GA four further metrics were recorded. These were the number of fitness gaining events (Gain); the number of events where no fitness gain was made (NoGain); the number of events where reproduction resulted in an existing genotype (Duplicates) and the CPU used in each reproductive event. The first three would indicate the effect of using the greedy algorithm to refine the P&P genotypes. The CPU metric on the other hand was significant as the land-block representation operates simply by looking up values within the DSS while the P&P GA is required to perform a series of more computationally intensive sorting operations in addition to the lookups.

It was hypothesised, (based on observation of runs during development), that while both representations would find acceptable solutions, the paths to those solutions would be different. The land-block representation GA would make a larger number of small magnitude gains, with fewer non-gaining events, while the P&P GA would proceed by larger magnitude steps, but with gains made more irregularly.

### 7.3 Initial Results and Discussion

Figure 3 illustrates the performance (FitMax and FitAvg) of each GA.




Figure 3. GA performance Graphs

Table 2 presents the mean value, over 50 runs, for the metrics outlines in Section 7.2). Counts of the All, Gain, NoGain and Duplicates were tabulated with the proportions of each class were also noted. The differences between the representations (Diff) for each metric were calculated and tested for significance (Sig). A Monte-Carlo sample-difference test [12] was used due to the skew in the distribution caused by the fixed upper limit of fitness values. The value of Sig is the probability the differences are not significant. The mean value for the count of CPU seconds was used to compare the computational effort expended.

Table 2. GA Performance

| Metric     | Land-block |     | P&P        |     | Diff   | Sig  |
|------------|------------|-----|------------|-----|--------|------|
| MaxFit     | £1,852,000 |     | £1,858,000 |     | £6,000 | 0.08 |
| AvgFit     | £1,851,000 |     | £1,846,000 |     | £5,000 | 0.10 |
| All        | 3654       |     | 1271       |     | 2383   | 0.00 |
| Gain       | 1405       | 38% | 722        | 57% | 683    | 0.00 |
| NoGain     | 1471       | 40% | 174        | 14% | 1297   | 0.00 |
| Duplicates | 778        | 21% | 375        | 30% | 403    | 0.00 |
| CPU(sec)   | 3319       |     | 3630       |     | -311   | 0.25 |

For the median value of maximum fitness both the P&P and land-block representations achieve 99% of the optimum value of £1,866,000. For both maximum and average population-fitness the performance differences between the two representations are not statistically significant.

The smaller number of reproductive events required to by the P&P GA to achieve these fitness values is statistically significant. Therefore as differences in average population-fitness at initialisation are insignificant the average fitness gains per reproductive cycle were as expected, with the P&P GA making fewer but larger gains (£850 to £1800).

Within the total event count there are further contrasts between the two representations with the P&P achieving fitness-gains more regularly than the land-block GA. The contrast is even more marked when comparing the NoGain event count.

While both representations make use of DSS suitability information to restrict initialisation and mutation, thus eliminating a large number of low-fitness allocations, the P&P GA also receives feedback from the greedy allocation algorithm. It uses this to repair the genotype structures and refine the target land use percentages as noted in Section 4.3. The use of the

feedback mechanism does, however, significantly increase the number of duplicate genotypes created when P&P target land use percentages are reduced to the percentage actually allocated. The number of duplicates is also increased by the use of the 5% granularity parameter for the P&P GA's target land use percentages.

Despite these differences in operation, when overall computational effort required is considered, there is no significant difference between the two representations. The reduced number of reproductive events for the P&P GA being bought at the cost of a balancing increase in computational effort.

## 8 Conclusions

In conclusion, while both representations find acceptable land use allocation solutions, and do so by equally efficient though differing routes, there are significant contrasts that can serve as guides to their use.

The P&P GA will be the preferred representation in cases where the application requires a substantial scaling up of the number of land blocks to be allocated. The caveat to this conclusion is, however, that the P&P representation can be expected to be sensitive to an increase in the number of land uses being considered as this will increase the average length of genotype.

The land-block representation has complementary features, being insensitive to the complexity of the optimum solution but adversely affected by significant increase in the numbers of land blocks. The land-block representation has the further advantage that it is more natural to use in the context of a multiple-objective search and optimisation, as it does not require the sorting function(s) associated with the P&P representation.

## Acknowledgements

This research is funded by SOAEFD grant-in-aid.

## References

- [1] Bramlette, M.F. and Bouchard, E.E. (1991). Genetic Algorithms in Parametric Design of Aircraft. In: Davis, L., (ed.) Handbook of Genetic Algorithms, 109-123. Von Nostrad Reinhold, New York.
- [2] Cartwright, H.M. and Harris, S.P., (1993). The Application of the Genetic Algorithm to Two-Dimensional strings: The Source Apportionment Problem. In: Forrest, S. (ed.) Proceedings of the Fifth International Conference on Genetic Algorithms, 416-423. Morgan Kaufmann, San Mateo, CA.
- [3] Cox, L.A., Davis, L. and Qiu, Y. (1991). Dynamic Anticipatory Routing in Circuit-Switched Telecoms Networks. In: Davis, L., (ed.) Handbook of Genetic Algorithms, 124-143. Von Nostrad Reinhold, New York.
- [4] Davis, L., (1991). A genetic algorithms tutorial. In: Davis, L., (ed.) Handbook of Genetic Algorithms, 1-101. Von Nostrad Reinhold, New York.
- [5] De Jong, K.A., (1990).. Introduction to the second special issue on genetic algorithms. Machine Learning 5(4):351-353.
- [6] De Jong, K.A., (1993). Genetic Algorithms are NOT function optimisers. In: Whitley, L.D. (ed.): Foundations of Genetic Algorithms 2, 5-17. Morgan Kaufmann, San Mateo, CA.

- [7] Eshelman, L.J., Caruana, R.A., and Schaffer, J.D. (1989). Biases in the crossover landscape. In: Schaffer, J.D. (ed.): Proceedings of the Third International Conference on Genetic Algorithms, 10-19. Morgan Kaufmann, San Mateo, CA.
- [8] Goldberg, D.E., (1989). Genetic Algorithms in Search, Optimisation and Machine Learning. Addison-Wesley, Reading MA.
- [9] Goldberg, D.E., Deb, K., Kargupta, H., and Harik, G., (1993). Rapid, Accurate Optimisation of Difficult Problems Using Fast Messy Genetic Algorithms. In: Forrest, S. (ed.) Proceedings of the Fifth International Conference on Genetic Algorithms, 56-64. Morgan Kaufmann, San Mateo, CA.
- [10] Matthews, K.B., Sibbald, A.R. and Craw, S (in Press). Implementation of a spatial decision support system for rural land use planning: integrating GIS and environmental models with search and optimisation algorithms. Computers and Electronics in Agriculture.
- [11] Michalewicz, Z., (1992). Genetic Algorithms + Data Structures = Evolution Programs. Springer, Berlin.
- [12] Noreen, E.W., (1989) Computer-intensive methods for testing hypotheses: an introduction. John Wiley & Sons, London
- [13] Samet, H., (1989). The design and analysis of spatial data structures. Addison Wesley, Reading, MA.
- [14] Syswerda, G., (1989). Uniform crossover in Genetic Algorithms. In: Schaffer, J.D. (ed.): Proceedings of the Third International Conference on Genetic Algorithms, 2-9. Morgan Kaufmann, San Mateo, CA.
- [15] Whitley, D., (1989). The GENITOR algorithm and selection pressure - why rank-based allocation of reproduction trials is best. In: Schaffer, J.D. (ed.): Proceedings of the Third International Conference on Genetic Algorithms, 133-140. Morgan Kaufmann, San Mateo, CA.