

Evaluating Multi-Objective Land Use Planning Tools Using Soft Systems Methods

K.B. Matthews, Susan Craw, S. Elder and A.R. Sibbald

Land Use Systems Group
Macaulay Land Use Research Institute
ABERDEEN, U.K.
AB15 8QH
Tel: +44 01224 318611
Fax: +44 01224 311556
k.matthews@mluri.sari.ac.uk
a.r.sibbald@mluri.sari.ac.uk

School of Computer and Mathematical Sciences
The Robert Gordon University
ABERDEEN, U.K.
AB25 1HG
Tel: +44 01224 262702
Fax: +44 01224 262727
s.craw@scms.rgu.ac.uk
s.elder@eee.rgu.ac.uk

Abstract

This paper presents a soft systems evaluation of two multi-objective genetic algorithms (mGAs) applied to land use planning. These mGAs search for a set of alternative solutions defining the structure of the relationship between two or more objectives. To investigate the usefulness of the mGAs for decision support the sets of solutions found by the mGAs are compared with allocations made by professional land managers, with a range of backgrounds, for a real world application. Conclusions are drawn on the differences between the solutions found by the land managers and the mGAs, in particular the constraints that limit the range of plans considered acceptable.

1 Introduction

Land use planning is a specialized domain of AI planning concerned with the optimization of spatio-temporal patterns of land use. The domain shares approaches with classical AI planning and scheduling, for example in the use of evolutionary computational methods for search and optimization. The planning tools presented here can, however, be usefully differentiated from AI planners by their search for target pattern(s) of land use to achieve objective(s) rather than planning the steps required, or the schedule of such steps, to proceed from the current to a target pattern of land use. In this regard they can be viewed as strategic or "marketing" level tools setting out possible options. AI planners and schedulers could of course inform the evaluation of the strategic planning tools' solutions determining the cost or feasibility of the transition given the constraint environment.

Even at a strategic level land use planning problems are complex with many spatially and temporally interdependent factors. This complexity is compounded when it is necessary to consider multiple, non-commensurable goals. This is particularly relevant as the context in which land management decisions are made is rapidly changing from one where production maximization is the single goal to one where land managers are expected to fulfill multiple, and frequently conflicting, financial, social and environmental objectives.

To assist land managers in their multi-objective, strategic land use planning tasks the Land Allocation Decision Support System (LADSS) is being developed (Matthews et. al 1999b).

Figure 1 shows the sub-components of the system. The geographic information system (GIS) provides the spatially referenced biophysical (climatic, soil and topographic) information and the facility to visualize the patterns of land use allocations as maps. The spatial biophysical information, with management and global parameters such as market prices, is used within the land use systems models to determine the suitability, productivity and profitability of individual parcels of land. The patterns of land use as a whole can be analyzed for their financial, social and environmental effects within the impact assessment component. The analytical core of LADSS is the iterative system, based on multi-objective genetic algorithms (mGAs), used to generate a set of land use allocations defining the structure of the relationship (usually a trade-off) between two or more objectives. The impact assessment component provides the evaluations needed to determine the relative utility or fitness of individual solutions. The allocations generated by the land use planning tools may be visualized within the GIS, have individual features queried or overridden or be subjected to further analysis with additional impact assessments.




Figure 1: Sub-components of LADSS

This paper compares, using soft systems methods, the performance of the two mGAs with that of land managers, on a test application, to qualitatively evaluate the systems' usefulness for decision support. After reviewing relevant related research in section 2 the features of the two mGAs within LADSS are outlined in section 3. The evaluations used are set out in section 4. The results for the test application are presented and discussed in section 5, and conclusions drawn in section 6.

2 Related Research

For mGAs the use of rank-based fitness assignment is so common that it may be viewed as one of the defining feature for the algorithms. Rank-based fitness assignment is necessary when considering two or more non-commensurable objectives, with the ranking most commonly based on the dominance relationship. A genotype is said to dominate another if superior in all evaluation dimensions. The mGAs within LADSS use the Pareto-ranking scheme as proposed by Fonseca and Fleming (1998) where the rank is equal to the count of the number of dominating individuals. Alternative methods based on normalization or

weighting of objectives have been used but the results of such methods depend heavily on their parameterization which may be difficult to determine a priori (Van Veldhuizen and Lamont 1998). The single solution from such methods also limits their usefulness as tools for the investigation of the relationships between objectives. The goal of the LADSS mGAs, is thus the search for non-dominated or Pareto-optimal solutions.

As the mGA population size is finite, it is usually the case that the mGA cannot find all Pareto-optimal solutions and the mGA thus has to search for a representative sub-set of the Pareto-optimal solutions that enables the structure of the relationship between objectives to be characterized. With two dimensions the Pareto-set forms a front that may be graphed. The most common approach to finding this representative subset has been the use of niche induction methods. These niching methods operate by progressively reducing the fitness of individuals (sharing) within a set distance (the niche size, a parameter of the GA). In mGAs such sharing is typically limited to solutions of the same rank and results in a selection bias in favor of non-dominated points spaced at a distance approximating the niche size. The appropriate setting of niche size has been noted as problematical (Zitzler, 1999). The domain in which the distance measurements are made may be genotypic (genetic similarity) or phenotypic (similarity of fitness evaluation). Arguments can be put forward for both. Genotypic sharing ensures genetic diversity and allows solutions with similar evaluations but differing genetics to coexist in the population. Phenotypic sharing ensures that solutions are evenly spread across the fitness function dimensions of interest to the decision maker (Horn et. al. 1994). For the mGAs within LADSS phenotypic sharing has been used to date because we are primarily interested in characterizing the form of the trade-off front.

Closely linked to the issue of niching is the mGA replacement strategy. Individual replacement or other elitist replacement strategies are superior for mGAs (Zitzler, 1999). This is because the mGA is searching for a set of co-adapted solutions, similar to a learning classifiers system's (LCS) search for co-adapted sets of rules (Valenzuela-Rendón and Uresti-Charre, 1997). With LCS it is recognized that the set of rules will only be found incrementally with individuals removed from the population only once they cease to be fit in the context of the population. The mGAs within LADSS use individual replacement.

The limitation of mGAs, as formulated with Pareto-rank based fitness and niching, particularly with individual replacement, is that they are computationally complex, because of the need to calculate distances between all members of the population. This has led to alternative mGA formulations using auxiliary populations to hold Pareto-optimal solutions with clustering used to reduce the auxiliary population size, while ensuring that the remaining individuals are representative of the Pareto-optimal solutions found (Zitzler, 1999). To date computational efficiency has not been a significant limitation within LADSS, particularly as the evaluation of individual solutions by the DSS is the main computational burden. Such methods may need to be adopted if significantly higher dimensional problems than the two dimensional problems considered to date are tackled.

3 LADSS mGAs

Two genotype representations have been used for the mGAs within LADSS (Matthews et al., 1999a). In the first representation the individual genes encode the land use for one land parcel. There is thus a direct, fixed-length mapping from gene to land parcel (Figure 2). This

is the land block (LB) representation. While this representation is both intuitive and requires no particular specialization of the mGAs, there was a concern that it would not scale up efficiently. To tackle this scaling issue the second representation was implemented where the mGA provides the parameters used to guide a greedy algorithm which makes the allocations to individual land parcels. The parameters required by the greedy algorithm are the target percentage of land to be allocated to a particular land use and the priority (order) in which land uses are allocated (Figure 3). This percentage and priority (P&P) representation scales with the number of land uses present, rather than the number of land parcels (usually at least an order of magnitude larger). The representation is, however, significantly more complex, with variable length genotypes, order-dependent interpretation of genes and messy elements such as under- and over-specification permitted (Goldberg et. al. 1993). This required the use of an enlarged operator set and the use of repair and post-evaluation processing of genotypes to remove functionally identical genotypes. Examples of this post-processing include the removal of genes that failed to result in the allocation of any parcels (the second wheat gene in Figure 3) and the combination of sequential pairs of genes allocating the same land use (the pair of forestry genes in Figure 3). For single objective optimization both representations found acceptable solutions but with the P&P GA using fewer learning cycles but the same total computational effort as the LB GA.




Figure 2: Land Block Representation

The two mGAs being evaluated share a common structure with each genotype representation supported by its own set of operators (Matthews et al., 2000). The populations of both mGAs are of fixed size, unstructured and have genotype uniqueness enforced. The enforcement of genotype uniqueness, as a precondition for insertion into the population, and before evaluation by the DSS, increases efficiency by avoiding unnecessary and computationally expensive fitness evaluations by the DSS. Following random initialization, both populations can be doped by adding genotypes derived from heuristics, the current pattern of land use or previous single objective GA runs. Doping assists the mGA performance as it introduces to the population solutions with close to optimal performance for one objective that may later be recombined with other members of the population to find intermediate solutions. The population size is set, based on niche size, using Fonseca and Fleming's (1998) formula.

Individual genotypes are translated into allocations within the DSS and the impact assessment component (Figure 1) provides the fitness assessments for each objective. These raw fitness values are translated into genotype rankings using the count of the number of dominating




Figure 3: P&P Representation

genotypes. Selection fitness is assigned using a linear normalization function, with genotypes sharing the same rank having their selection fitness values averaged. Fitness values are shared between genotypes of the same rank using the triangular sharing function proposed by Horn et al. (1994).

The operator choice, selection, reproduction and replacement cycle is organized as follows. Operators are independently applied with their probability of use adapted over time (Davis 1991). The first parent is selected using roulette wheel selection; if a second parent is required the selection fitness values are biased by reducing the selection fitness of genotypes outwith a defined mating distance (Fonseca and Fleming 1998). The offspring from the operator is/are then checked for uniqueness against the existing population, evaluated (and in the case of P&P genotypes possibly modified by the feedback from the DSS), and inserted individually into the population. At each insertion the population is ranked, fitness values assigned and shared. The lowest fitness genotype is then eliminated.

Both mGAs employ a uniform crossover operator, with unequal length P&P genotypes exchanging genes only between the initial segment common to both parents. The operator set for the LB mGA is completed by a single mutation operator that changes the current land use for a gene to another suitable land use (Matthews et.al. 2000). The P&P mGA has two further binary operators, splice and order prioritized crossover. The first simply concatenates two existing genotypes to produce a single offspring. The second selects a subset of the genes using a crossover mask and then reorders those genes based on the order in which they appear in the other parent. The P&P mGA requires three mutation operators. Type-mutation changes the class of the gene and thus the land use that it causes to be allocated. Non-uniform mutation changes the target land use percentage. Pair-swap exchanges the genes at two loci thus altering the priority with which land uses are allocated. The final two P&P operators

are insert-gene and delete-gene which increase or decrease the length of a P&P genotype by adding a new randomly generated gene or deleting an existing gene selected at random.

4 Evaluation

From previous analysis of the performance of the LB and P&P mGAs in a two-dimensional problem (Matthews et. al. 2000) it was found that, in terms of the individual objectives, the P&P mGA marginally outperformed the LB mGA. This difference was in the consistency with which the P&P mGA found the individual objective optima. For the measures of solution quality both mGAs were equally efficient in finding non-dominated solutions (98 and 99% of the terminal populations for P&P and LB respectively). The evenness of coverage for the LB mGA was superior to that of P&P because of the use of 5% granularity in the specification of target land use percentages in the P&P genes. This meant that the P&P mGA was less able to generate solutions evenly spread across the front. The superior performance of the LB mGA in coverage and dominance was again in the consistency with which it found intermediate solutions. The P&P mGA terminated in the majority of runs when it exceeded the consecutive no-gain limit of twenty reproductive events and generated much larger numbers of duplicate genotypes (38% compared to 1% for LB). It was therefore concluded that the population size was not large enough to preserve genetic diversity given the smaller number of genes used in the P&P representation. Both mGAs, however, provided useful means of characterizing the structure of the trade-off front between the two objectives.

Beyond the hard evaluation of mGAs it has been seen that the evaluation of the functionality, performance and usefulness of decision-supporting software may be investigated using soft systems methods (Van Beek, 1995). Soft systems appraisal methods are workshop based with delegates chosen to represent a range of differing perspectives and priorities. In the land management domain this would include landowners, non-governmental interest groups, rural investment institutions (banks) and academics. Typically, soft-systems workshops use facilitated sub-groups to produce qualitative analyses with these analyses compared in plenary sessions (van Beek and Nunn 1995).

The soft-systems evaluation of the two mGAs' performance was designed to capture both a qualitative assessment of the usefulness of the tools and to compare for a test application the solutions found by the mGAs with those produced by experienced land managers or systems analysts, (with differing land management priorities).

The test application is a farm in Lanarkshire, Scotland and was the same as that used in previous testing of the mGA performance using hard metrics (Matthews et al., 2000). The workshop delegates were first asked individually to design a land use allocation pattern that would maximize financial returns and land use diversity. The financial metric was the gross margin (income minus inputs costs but not including capital or labour costs) expressed as a net present value (NPV) over 60 years (Boehlje and Eidman, 1984). The landscape diversity index used was the Shannon-Wiener (SW) index that is maximized when all potential land uses are present in equal proportions (Forman and Godron, 1992). These two objectives are known to be antagonistic as increasing areas of less financially productive land uses such as forestry increase the SW index while reducing the NPV. Given the known trade-off between the two objectives the delegates were then asked to produce what was termed a "workable compromise". This gave each delegate the scope to balance the objectives given their varying

perspectives, from conservation to estate management. The aim was thus to produce a population of alternative allocations.

The delegates based their allocation design on an information pack containing maps, photographs and tables of information on the climate, soils, topography and current land uses of the farm. The information pack was designed to give sufficient detail to allow informed decisions to be made without prejudicing the range of allocations produced. The information provided was broadly similar in nature to that which would be available to a land management consultant.

Following the completion of the individual allocations, the delegates were divided into two sub-groups (see Table 1) and asked to produce a compromise solution from the sub-group as a whole. A member of the research team facilitated the process of deriving the group allocation, with each member of the group presenting their individual plans and the group working together to answer the following questions. Is the plan workable a whole, are there parts of the plan that must be kept/dropped and are there elements that can be added to improve the plan? Following the presentations the groups were asked to agree a plan by first defining the elements that are fixed and non-negotiable, second consensus allocations, third land uses/locations that must not occur and finally areas where any land use would be acceptable.

Subgroup 1	Subgroup 2
BA1 - bank adviser	SA2 - systems analyst
AG1 - agriculturalist	AG2 - agriculturalist
B1 - biologist	C2 - conservationist
E1 - estate manager	E2 - estate manger
F1 - farm manager	F2 - farm manager

Table 1: Delegates by sub-group

In order to ensure that the allocations produced by the delegates and the sub-groups could be analyzed within LADSS it was necessary to propose a series of simplifying assumptions. These were: 22

1. The land allocation is defined per existing land parcel from the range of possible land uses. This assumption was maintained but for the test application only a subset (5) of the possible (10) land uses were considered of practical value (trees were allocated under a general classification of broad leaved and conifer species rather than by individual species). This restricted set of land uses (arable, cattle, sheep, broad leaved and conifer trees) was also imposed on the mGAs to simplify the process of comparing results. The potential for diversification was noted and several alternative land use strategies were therefore disregarded.
2. No changes to the existing pattern of field boundaries. This was accepted but noted as limiting for certain delegates' plans.
3. No land may be bought or sold. One solution was proposed (G1-2 in following section) with surplus land assumed to be rented to other land managers.
4. The existing land uses do not limit future potential. This was accepted but all delegates

went further and retained all existing woodland thus fixing 5% of the farm as common to all allocations.

5. Capital and infrastructure are not limiting. While accepted, this was highlighted as one of the key constraints on real-world land use change.

Over the course of the workshop it also became clear that in discussing allocations, additional metrics were being used. These included total stock numbers, and summaries of the arable and forestry land use by percentage of the total enterprise area. These have been adopted as secondary metrics for comparing delegates' land allocations.

5 Soft Systems Results and Discussion

The eight delegates produced ten individual allocations and three group allocations. The current pattern of land use was also analyzed. The graph in Figure 4 shows the location of the delegates' plans in the search space, defined by the two fitness functions. The same figure also illustrates the Pareto-sets of plans found by the P&P and LB mGAs. Table 2 presents for each delegate allocation the NPV value, the SW index, the numbers of animals and the area of broad leaved and conifer trees and arable crops.

The upper limit on SW values of 1.6 is a 20% by area allocation to each of the possible land uses (the upper dotted line in Figure 4). The lower limit of zero occurs when a monoculture is imposed. It is clear from Figure 4 that the delegates' solutions occupy only 50% of the possible range of SW values. The first reason for this is that all delegates imposed a precondition that all existing woodland would be preserved. This immediately removes the possibility of a cattle monoculture, the financial optimum found in previous analysis(Matthews et. al, 1999a) and sets a lower bound on diversity of 0.25 (the lower dotted line in Figure 4) and an upper bound on financial returns of £5.25M pounds. (This precondition was also imposed on the mGA search). The second and more important increase in the lower bound for SW values is that, while it is possible to run single species livestock farms, there are good animal welfare reasons for having both sheep and cattle present within a single enterprise. The area of land devoted to sheep is usually less than that for cattle but must be sufficiently large that it can link with the rotation of cattle land. A sheep-cattle mixed livestock system with existing trees preserved raises the lower diversity limit to approximately 0.8 (the middle dotted line in Figure 4).

Within these SW bounds the delegates proposed allocations across a range of financial returns from £1.59M to £4.00M. The distribution of these allocations can be seen to roughly form a trade-off front similar in shape to that found by the mGAs with, in most cases, marginally (10-15%) "poorer" performance than the mGA. This difference in financial performance is the result of factors not explicitly taken into account by the mGA search. An example of one such factor is the delegates blocking together of fields into management units, for example keeping all the sheep fields contiguous. There are good practical management reasons for doing this but it does mean that the expert allocations may not necessarily be optimized.

The three individual allocations furthest from the front (B1, BA1 and E1-1) are useful in indicating that there are possible solutions throughout the search space, and that solutions to the test problem are not necessarily clustered close to the Pareto-front. For both solutions

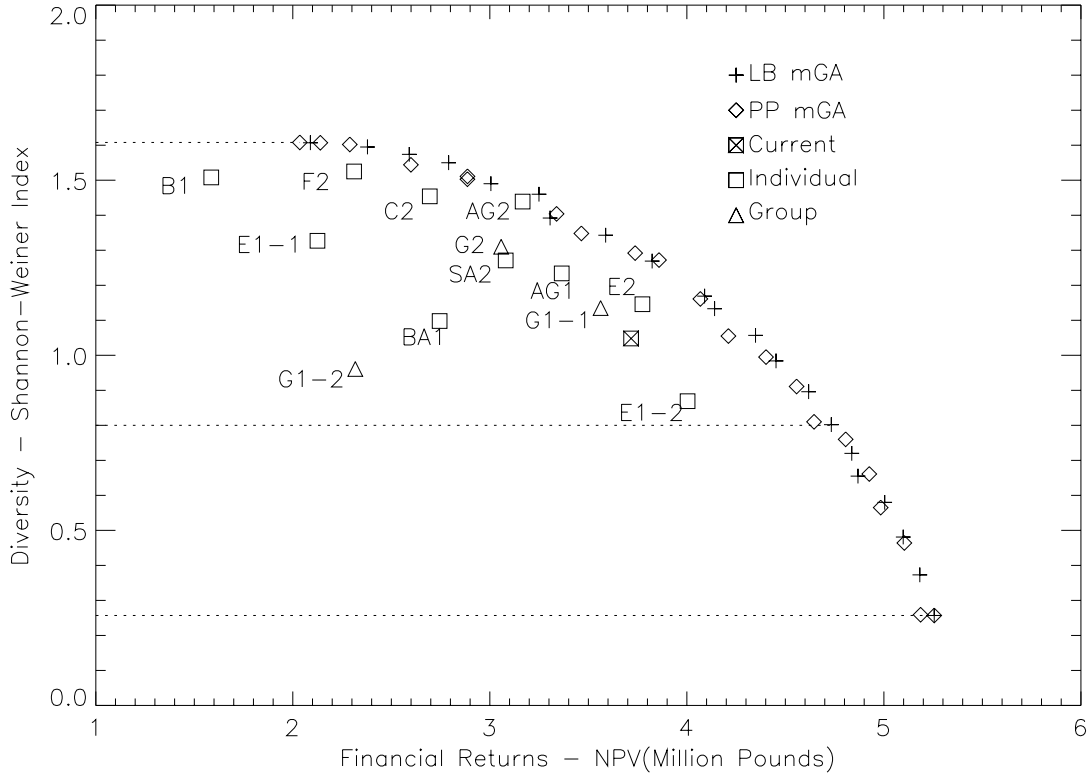


Figure 4: mGA and Land Manager Allocations

Delegate Code	NPV (£M)	SW Index	Sheep (No.)	Cattle (No.)	Broadleaf (Ha)	Conifer (Ha)	Arable (Ha)
Current	3.71	1.048	1223	348	23.3	0	15.2
E1-2	4.00	0.869	1187	401	21.8	0	0
E2	3.77	1.146	1016	355	26.8	0	32.6
G1-1	3.56	1.135	967	347	49.0	0	14.4
AG1	3.36	1.234	544	329	78.5	0	33.9
AG2	3.16	1.439	597	307	42.9	40.7	33.9
SA2	3.08	1.271	1098	251	36.4	0	45.1
G2	3.05	1.31	802	271	76.0	0	38.6
BA1	2.74	1.098	1150	246	103.0	0	0
C2	2.69	1.454	591	255	54.9	21.7	31.5
F2	2.31	1.525	898	187	76.1	36.5	30.6
G1-2	2.31	0.961	1982	137	71.4	0	0
E1-1	2.12	1.327	1651	133	57.7	34.3	8.3
B1	1.58	1.508	768	110	113.6	47.9	29.7

Table 2: Delegate Allocations (ordered by NPV)

B1 and BA1 the solutions are financially sub-optimal due to the large extent of the low-value woodland areas (113 and 103 ha. out of a total area of 300 ha.). For E1-1 the reason for sub-optimality was the failure to allocate the arable land uses to land parcels defined as suitable by LADSS. While the information pack did provide the information on which suitability could be judged, it did not provide either the criteria used by LADSS or maps of the land blocks suitable for particular land uses. This was probably a fault in the analysis as it confuses the ability of the land managers to come up with compromise land allocations with their ability to determine the suitability of land for particular land uses. It was decided that, to minimize the possible impact of misallocation, the suitability criteria would be relaxed for arable crops (this can be rationalized as possible management interventions). While this eliminated almost all of the illegal arable allocations from the delegates (those in E-1-1 being exceptional in this regard), it also meant that the allocations found by the mGAs changed in character with much larger areas of arable crops present. The need for care in experimental design even in such soft systems analysis is apparent.

The sub-group G1-2 allocation is included in the results for the purpose of illustrating how carefully the mGA metrics have to be chosen and the degree of multi-dimensionality in land management issues. The G1-2 allocation proposed the inclusion of available labour as an additional precondition. In this case the allocation assumed a single, full-time labour unit available with a pattern of land use dominated by sheep, but with some forestry and the remainder of the land rented to other farmers for seasonal grazing. The financial analysis of such a system depends not only on the gross margins but also on the ratio of input costs (including labour and machinery) to output revenue. The system proposed had very low input costs and thus could be much closer to the financial optimum than indicated by the metric used for the current analysis. There are also potential environmental benefits from the less intensive farming regime but social costs because of the reduced levels of local employment provided.

6 Conclusions

Two mGAs with a shared structure differentiated by genotype representation and operator sets have been proposed. Previously conventional hard metrics had been used to evaluate their relative performance. The conclusion of this analysis was that while both mGAs found acceptable Pareto-sets the P&P mGA could be hampered by conflicts in its parameterization.

To further investigate the mGAs performance workshop-based soft systems methods were used to collect allocations made by land management specialists that could be compared with the Pareto-sets found by the mGAs. The comparison revealed that the practitioners operated within an agreed set of constraints that limited the range of allocations considered, but that, within those limits, solutions were found across the search space, and close to the Pareto-front defined by the mGAs. Practical management concerns, such as the desire for land parcels of some land uses to be spatially contiguous, was hypothesized as the most likely reason for differences between the practitioner allocations and those of the mGAs. The utility of the mGAs would be improved by the use of spatial contiguity information (provided by the GIS), either as a constraint or as an explicit optimization goal. The allocations found by the mGAs were, however, agreed by the land managers to be capable of forming the basis of management plans with operator-applied modifications to individual land parcels to ensure real world practicality.

Finally the soft systems analysis also provided a wide range of qualitative evaluations for both the mGAs and the DSS. These insights suggested improvements to the range of analyses the DSS should provide, the metrics used by land managers in planning/comparing land allocations, heuristics that could be added as default allocation strategies, and the key constraints required to ensure that the allocations found by the mGAs are workable. The workshop also provided anecdotal backing for the view that land management professionals faced with complex multi-objective planning problems want interactive decision support tools where a range of options can be examined and conclusions drawn on the trade-offs in costs and benefits. A wide choice of evaluation metrics and constraints that can be customized to the specific conditions of a particular land management unit is essential. The combination of hard metrics backed by soft systems-based analysis has thus proved highly effective both in evaluating the performance of the mGAs and in suggesting improvements to the range of analyses supported by the DSS.

Acknowledgments

This research is funded by grant-in-aid from the Scottish Executive Rural Affairs Department. The authors also acknowledge the assistance provided by Scottish Enterprise Grampian.

References

- Boehlje, M.D. and Eidman V.R. (1984) Farm Management. Wiley and Sons, NY.
- Davis, L. (1991) Handbook of Genetic Algorithms. Von Nostrand, Reinhold, NY.
- Fonseca, C.M. and Fleming, P.J. (1998) Multi-objective Optimisation and Multiple Constraint handling with Evolutionary Algorithms -Part 1: A unified formulation. *IEEE Transactions on systems Man and Cybernetics - Part A: Systems and Humans*. Vol. 28, No 1, pp. 26-37.
- Forman, R.T.T., and Godron, M., (1986) Landscape Ecology. Wiley and Sons, NY.
- Goldberg, D.E., Deb, K., Karagupta, H., and Harik, G. (1993) Rapid, Accurate Optimisation of Difficult Problems Using Fast Messy Genetic Algorithms. In: S. Forrest (ed.) *Proceedings of the Fifth International Conference on Genetic Algorithms*, 56-64, Morgan Kaufmann, San Mateo, CA.
- Horn, J., Nafpliotis, N. and Goldberg, D.E. (1994) A Niche Pareto Genetic Algorithm for Multiobjective Optimisation. In: Zbigniew Michalewicz (ed.) *Proceedings of the First IEEE Conference on Evolutionary Computation*, 82-87, Piscataway NJ: IEEE Service Centre.
- Matthews, K.B., Craw, S., MacKenzie, I., Elder, S., and Sibbald, A.R. (1999a) Applying Genetic Algorithms to Land Use Planning. In: G. Petley, A. Coddington and R. Aylet (eds.) *Proceedings of the 18th Workshop of the UK Planning and Scheduling Special Interest Group*, 109-115, University of Salford.
- Matthews, K.B., Sibbald, A.R. and Craw, Susan, (1999b). Implementation of a spatial decision support system for rural land use planning: integrating geographical information system and environmental models with search and optimisation algorithms. *Computers and Electronics in Agriculture* (23) 9-26.
- Matthews, K.B., Craw, S. MacKenzie, I., Elder, S., and Sibbald, A.R. (2000) Applying Genetic Algorithms to Multiobjective Land Use Planning. In: D. Whitley (ed.) *Proceedings of*

the Genetic and Evolutionary Computation Conference (GECCO 2000), 613-620, Las Vegas, USA.

Valenzuela-Rendón, M., and Uresti-Charre E. (1997) A Non-Generational Genetic Algorithm for Multiobjective Optimisation. In: Thomas Bäck (ed.) *Proceedings of the Seventh International Conference on Genetic Algorithms*, 658-665, Morgan Kaufmann, San Mateo, CA.

Van Beek, P.G.H. (1995) A Review of Decision Support Tools (DSTs). Macaulay Land Use Research Institute, 26 pages.

Van Beek, P.G.H. and Nunn, J. (1995) Introduction to Rapid Appraisals. SyTREC Pty Ltd.

Van Veldhuizen, D.A. and Lamont, G.B. (1998) Multiobjective Evolutionary Algorithm Research: A History and Analysis. Technical Report TR-98-03, Department of Electrical and Computer Engineering, Air Force Institute of Technology, 105 pages.

Zitzler, E., Deb, K., and Thiele, L. (1999) Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. TIK-Report No. 70. Institute für Technische Informatik and Kommunikationsnetze, ETH Zürich, 19 pages.