Using soft-systems methods to evaluate the outputs from multi-objective land use planning tools

K.B. Matthews[a], K. Buchan[a], A.R. Sibbald[a], and Susan Craw[b]

[a]*The Macaulay Land Use Research Institute, Aberdeen, AB15 8QH, U.K. (k.matthews@macaulay.ac.uk)*

[b]*School of Computing, The Robert Gordon University, Aberdeen, AB25 1HG, U.K.*

Abstract: Land managers are increasingly faced with complex decisions requiring the consideration of trade-offs between multiple, non-commensurable objectives. Such decisions have profound effects on the financial, social and environmental sustainability of land use systems. One approach to assisting land managers with these decisions has been the development of computer-based decision support systems (DSS). While such systems are demonstrably able to make analyses of multi-objective land-use planning problems, are the answers they produce relevant and useful to practitioners? This paper reports on a workshop-based, soft-systems analysis of outputs from a spatial, multi-objective land-use planning tool. The paper outlines the approach taken in developing the decision support system, focusing on the land-use planning tools. These tools use multi-objective genetic algorithms to define the structure of the trade-off between objectives. The paper then details the soft-systems-based evaluation strategy. Land managers and other professionals from a range of backgrounds were asked to devise individual "best compromise" plans, balancing financial and landscape diversity goals, for a farm in upland Scotland. Sub-groups of land managers were then set the task of agreeing on a plan between the members of the group. This process used the soft-systems methods of facilitated discussion and reporting back from sub-groups. The land managers' and sub-groups' plans were analysed with the DSS tools and the results compared with outputs from the land-use planning tools. From this analysis and the qualitative responses within the workshop it was possible to conclude that the land-use planning tools provided a useful means of exploring the patterns of land use that could be adopted for a land management unit. The process identified a number of assumptions made by land managers that could usefully be incorporated into the operation of the DSS. The use of soft-systems based analysis of land-use planning tool outputs is recommended, not only for evaluating the performance of the tools, but also for ensuring that the DSS is answering a correctly formulated problem.

Keywords: Land-use planning, multi-objective, genetic-algorithms, soft-systems

1 Introduction

Land managers are increasingly faced by the need to achieve production and financial goals within tighter social and environmental constraints. One response has been the development of computer-based decision-support systems (DSS). These DSS seek to assist land managers in exploring the options for, and impacts of, changes to their patterns of land use [Matthews et al., 1999b]. Particularly significant components of these DSS are multi-objective land-use planning tools. The integration of these tools with geographic information systems and simulation models provides a flexible analytical framework where the trade-off between financial, social and environmental impacts can be evaluated. How useful are these tools for real-world planning problems? This paper reports on a soft-systems based evaluation of the multi-objective land-use planning tools within a strategic land-use DSS.

2 Related Work

2.1 Multi-objective Planning

For multi-objective land-use planning methods the key distinction is between *a priori* and *a posteriori* approaches [van Veldhuizen and Lamont, 2000]. With *a priori* the strategy is *decide-then-search*, with the decision-maker defining a weighting or ordering scheme for the objectives from which a solution is generated. This strategy uses the scalarisation or ordering approaches seen in *multi-criteria decision making* (MCDM) for site selection [Carver, 1991] or indicative zoning [Beedasy and Whyatt, 1999]. Solutions found by *a priori* methods are known to be sensitive to the weightings and orderings employed [Fonseca and Fleming, 1995]. This means that in certain situations, particularly where there is conflict over a decision, it may be impossible to agree on these values. *A priori* methods may also be employed corruptly to justify a desired solution. In *a posteriori* strategies the approach is *search-then-decide*, with the decision-maker presented with a range of alternatives, defining the trade-off between objectives [Matthews et al., 2000].

2.2 Pareto-optimality

It is intuitive that for many land-use planning problems no Utopian solution will exist, where all objectives are simultaneously optimal, (Figure 1). The nature of "optimality" for the case shown in Figure 1 can be formalised using the *dominance* relation between alternative solutions [Goldberg, 1989]. One alternative *dominates* another, and is therefore preferable, only if superior in pairwise comparison, for at least one objective, and has equal or better performance for the other objectives. Alternatives that are not dominated by any other within the space defined by the objectives are members of the set of *Pareto-optimal* solutions [van Veldhuizen and Lamont, 2000]. These alternatives make the best possible compromise between objectives. To characterise the trade-off between objectives the land-use planning tool should find a set of non-dominated solutions (the PO-set) encompassing the range of fitness values that are Pareto-optimal and evenly spread across that range [Matthews et al., 2000].

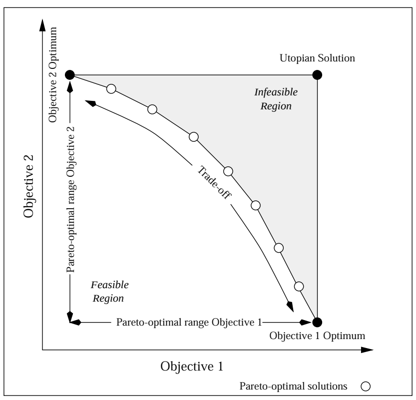

Figure 1. Two-objective maximisation example.

2.3 Multi-objective Genetic Algorithms

One successful approach to finding sets of non-dominated land-use plans defining the trade-off between objectives is to use multi-objective genetic algorithms (mGAs), [Matthews et al., 2000; Ducheyne et al., 2001]. mGAs are an extension to a class of search and optimisation algorithms based on the mechanics of natural selection [Goldberg, 1989]. GA's maintain a *population* of alternative solutions from which individuals *(termed genotypes)* are selected as *parents*. Components *(genes)* from the parents are recombined by *operators* to form new *offspring* solutions. Selection is biased in favour of those genotypes evaluated as *fitter* by so called *fitness functions*. Fitness will be application dependent and evaluated by the DSS. For mGAs, fitness is typically rank-based with Pareto-optimal solutions having the highest rank. The least fit member(s) of the population are replaced by offspring that are fitter.

mGAs have been shown to be robust and efficient algorithms for searching large, complex and little-understood search spaces such as those of multi-objective land use planning. mGAs are particularly effective as they allow, within a single run, the characterisation of the trade-off using the population. Limitations of mGAs include their stochasticity, need for careful parametrisation and the need to represent the problem in such a way that the GA operators can be effective.

Two genotype representations have been used for the mGAs [Matthews et al., 1999a]. In the land block (LB) representation individual genes encode the land use for one land parcel. There was a concern that with large numbers of blocks the LB representation would be computationally impractical. A second representation was therefore proposed where genes hold two parameters, the target percentage to be allocated and the priority for each land use. These parameters are used by a *greedy algorithm* [Goldberg, 1989] to iteratively allocate land blocks starting with those having the best performance for the highest priority land use. Allocation continues until either the target land use percentage is exceeded or no land blocks remain to be allocated. The computation required by the percentage and priority (P&P) representation is proportional to the number of land uses present, rather than the number of land parcels Matthews [2001]. For both representations the mGA may only allocate a use to a block if that use is bio-physically possible, as determined by the rules within the DSS.

There is a range of hard metrics and approaches for comparing the performance of alternative mGAs

[Zitzler and Thiele, 1998]. These metrics do not, however, address the usefulness of the mGA approach, in particular how the performance of the mGA compares to that of experienced land-managers. The remainder of this paper presents the analysis of a soft-systems evaluation of the land-use planning mGAs.

3 Materials and Methods

Soft-systems methods have been used to evaluate the required functionality for DSS [van Beek, 1995]. Soft-systems appraisal methods are workshop-based with delegates chosen to represent a range of differing perspectives [van Beek and Nunn, 1995]. In this study the workshop delegates included land-managers, interest groups, banks and academics. Typically, soft-systems workshops use facilitated sub-groups (SG's) to produce qualitative analyses with these analyses compared in plenary sessions.

The workshop delegates were first asked to design individually a pattern of land use for a farm in Lanarkshire, Scotland, previously used in the development and testing of the mGAs. Two goals were stated, to maximise financial returns and land use diversity. The metric for the financial goal was the farm gross margin (income minus input costs excluding capital and labour) expressed as a net present value (NPV) over 60 years [Boehlje and Eidman, 1984]. The land use diversity was measured using the Shannon-Wiener (SW) index that is maximised when all potential land uses are present in equal proportions [Forman and Godron, 1986]. These two objectives were 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 asked to produce the "best compromise". This gave each delegate the scope to balance the objectives given their varying perspectives. The aim was thus to generate a set of alternative land-use allocations.

To assist in their allocation design each delegate received an information pack containing 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 individual allocations, the delegates were divided into two SGs, Table 1 and asked to produce a compromise solution from the SG as a whole.

Table 1. Delegates by Sub-group

Sub-group 1	Sub-group 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

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 as a whole? Are there parts of the plan that must be kept/dropped? Are there elements that can be added to improve the plan? Following the presentations of the individual plans, the groups were asked to agree on a plan by firstly defining the elements that are fixed and non-negotiable, secondly consensus allocations, thirdly patterns that must not occur and finally areas where any land use would be acceptable.

In order to ensure that the allocations produced by the delegates and the SGs could be analysed within the DSS it was necessary to propose a series of simplifying assumptions. These were:

1. The land allocation is defined per existing land parcel from the range of possible land uses. This assumption was maintained but only a subset (5) of the possible (10) land use were considered of practical value; arable, upland-sheep and suckler-cattle with broad-leaved and coniferous trees allocated rather than individual species. The restricted set of land uses was also imposed on the mGAs to simplify the process of comparing results. The potential for diversification out of farming was noted.

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.

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 9% 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.

4 RESULTS AND DISCUSSION

4.1 mGA and Delegate Allocations

The eight delegates produced ten individual allocations and two group allocations. The current pattern of land use was also analysed. Figure 2 shows the location of the individual delegates' plans in the search space, defined by the two fitness functions. The same figure also illustrates the PO-sets found by the P&P and LB mGAs.

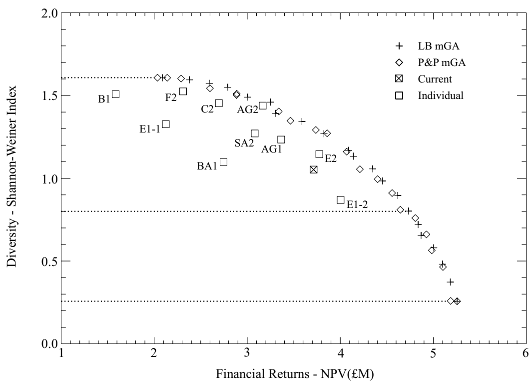

Figure 2. mGA and Land Manager Allocations

The upper limit on SW values is 1.6 and represents a 20% allocation to each of the possible land uses (the upper dotted line in Figure 2). The lower limit of zero occurs when a mono-culture is imposed. It is clear from Figure 2 that the delegates' solutions occupy only part of the possible range of SW values. The delegates first imposed a precondition that all existing woodland would be preserved. This immediately removes the possibility of a cattle monoculture (the financial optimum) and sets a lower bound on diversity of 0.25 (the lower dotted line in Figure 2) and an upper bound on financial returns of £5.25M. This precondition was also imposed on the mGA search. 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 farm. The area of land devoted to sheep is usually less that that for cattle but must be sufficiently large that it can link in a rotation with cattle. 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 2).

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 front similar in shape to the PO-set found by the mGAs with, in most cases, marginally "poorer" performance than the mGA. The mean magnitude of this difference is -26% for financial return and -13% for land-use diversity. It was hypothesised that the financial performance of the delegates' allocations was reduced by their assembling of fields into higher-level 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 allocations may not necessarily be optimal for the fitness-functions imposed.

To test this hypothesis the P&P algorithm was altered as follows. After an initial land-block has been allocated to a land-use by the greedy algorithm the list of remaining land-blocks is sorted by proximity to the initial land-block. Allocation then continues on a closest-first basis. Figure 3 plots the PO-sets found by the modified *proximity-P&P* mGA against the delegate allocations.

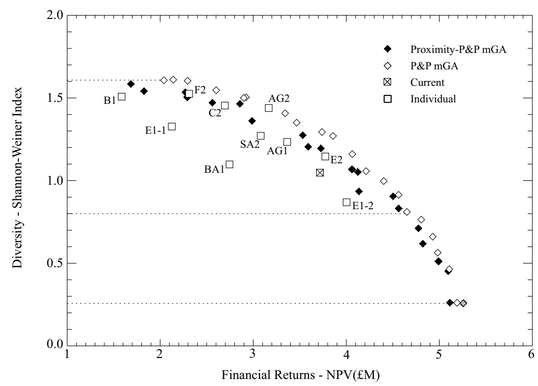

Figure 3. Proximity-P&P vs. land managers

It is clear that assembling blocks into management units reduced the NPV values for allocations with the same diversity. The difference in NPV between the P&P and proximity-P&P PO-sets thus becomes more significant as the diversity increases. More of the solutions proposed by the delegates lie close to the proximity-P&P PO-set, with the mean difference reduced to -13% for finance and -5% for diversity.

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 PO-set. For both solutions 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 the DSS.

4.2 Sub-group Allocations

Figure 4 shows the SG allocations (G1 and G2) and the allocations of the delegates who contributed to each. The figures also show the centre-of-gravity (CoG), an unweighted average of the coordinates from the delegates' individual allocations. The CoG is the expected location of a group solution based on an equal compromise between the delegates.

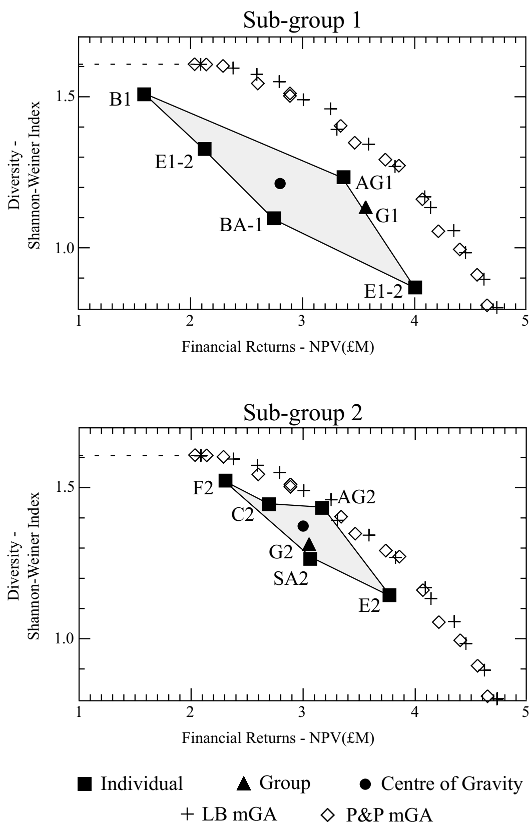

Figure 4. Sub-group allocations

It is clear that for both SGs there has been compromise between the delegates, but also that that the compromise has not been on an equal basis. By calculating weighted averages of individual allocations' coordinates it was possible to explore the influence of individuals on the SG allocation. Weights were set using a gradient-descent method, minimising the difference between the weighted average and the coordinates of the SG allocation. Initially the weights were biased to be non-negative and sum to one but Monte-Carlo testing found that, since G1 lies outside the polygon with vertices at the delegate allocations, it is necessary to allow the sum of weights to marginally exceed one. Graphing the weights from 100 gradient-descents revealed that there were several combinations of weights that would result in the SG allocation. To reflect this uncertainty the probability of each delegate being the most influential (PMI in Table 2) was calculated. For SG-1, E1-2 was the most influential with AG1 and BA1 second with nearly equal probabilities. For SG-1 two solutions were never the most influential. For SG-2 only one delegate was never the most influential. For the others there is a ranking but without a strong bias. These PMI values match the perceptions of group interactions by the SG facilitators.

In addition to the PMI metric the mean value of the weights (Mean WT) per individual was calculated, Table 2. The mean value of weights is not an ideal measure but does allow limited inferences to be drawn on the influence of delegates that are never the most influential. For SG1, it is clear that E1-1 is more influential than B1 and for SG2 that F2, while never the most influential, is not on average greatly less influential than AG2. For both the PMI and Mean WT metrics the greatest influence is not exerted by the delegate solution closest in the search space to the SG allocation.

Table 2. Delegate influence

Sub-group	Delegate	PMI	Mean WT
1	BA1	20	0.22
1	AG1	16	0.21
1	B1	0	0.04
1	E1-1	0	0.12
1	E1-2	64	0.48
2	SA2	26	0.19
2	AG2	22	0.17
2	C2	18	0.19
2	E2	34	0.29
2	F2	0	0.14

While this analysis of the compromise process is limited it does highlight the complexities of group-based decision making. The delegates were clearly willing to compromise, though the degree, had real incomes or environmental impacts been at stake, is less certain. The effectiveness of the SG compromises is questionable. While distance to the mGA PO-set is a crude measure of quality, it does illustrate that the SG solutions do not improve on the allocations made by the best individuals. This is particularly noticeable for SG2 where the near-equal influence of the delegates results in a solution further from the PO-set than all but one of the individual allocations. For SG1, the strong influence of E1-2 is evident in improving the financial performance of G1. In both cases the compromise process, unlike the mGA, is losing the best elements of the individual solutions not recombining them together to make better solutions.

5 CONCLUSIONS

Workshop-based soft systems methods were used to collect allocations made by land management specialists that could be compared with the PO-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 in the majority of cases, close to the PO-set found by the mGAs. Practical management concerns, such as the desire for land-blocks of some land uses to be spatially contiguous, was shown as a significant reason for differences between the practitioner allocations and those of the mGAs. The allocations found by the mGAs were, however, agreed by the land managers to be capable of forming the basis of management plans with modifications to individual land-blocks to ensure real-world practicality.

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 and 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.

ACKNOWLEDGEMENTS

The authors greatfully acknowledge the funding support of the Scottish Executive Environment and Rural Affairs Department and Scottish Enterprise Grampian and the significant contribution of the workshop delegates to this research.

REFERENCES

Beedasy, J. and D. Whyatt. Diverting the tourists: a spatial decision-support system for tourism planning on a developing island. *JAG*, 1(3–4):163–174, 1999.

Boehlje, M. and V. Eidman. *Farm Management*. Wiley and Sons, New York, 1984.

Carver, S. J. Integrating multi-criteria evaluation with geographical information systems. *International Journal of Geographical Information Systems*, 5(3):361–367, 1991.

Ducheyne, E., R. de Wulf, and B. de Baets. Bi-objective genetic algorithm for forest management: a comparative study. In Goodman, E., editor, *Proceedings of the 2001 Genetic and Evolutionary Computation Conference (GECCO 2001)*, pages 63–66. Morgan Kaufmann, San Francisco, July 2001.

Fonseca, C. M. and P. J. Fleming. An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation*, 3(1): 1–16, Spring 1995.

Forman, R. and M. Godron. *Landscape Ecology*. Wiley and Sons, New York, 1986.

Goldberg, D. *Genetic Algorithms in Search, Optimisation and Machine Learning*. Addison-Wesley, Reading, MA, 1989.

Matthews, K. *Applying genetic algorithms to multi-objective land-use planning*. PhD thesis, The Robert Gordon University, Aberdeen, U.K., May 2001.

Matthews, K., S. Craw, S. Elder, A. Sibbald, and I. MacKenzie. Applying genetic algorithms to multiobjective land use planning. In Whitley, D. et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2000)*, pages 613–620. Morgan Kaufmann, San Francisco, 2000.

Matthews, K., S. Craw, I. MacKenzie, S. Elder, and A. Sibbald. Applying genetic algorithms to land use planning. In Petley, G. et al., editors, *Proceedings of the 18th Workshop of the UK Planing and Scheduling Special Interest Group*, pages 109–115, University of Salford, 1999a. UK Planing and Scheduling Special Interest Group.

Matthews, K., A. Sibbald, and S. Craw. 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*, 23:9–26, 1999b.

van Beek, P. A review of decision support tools (DSTs). Technical report, Macaulay Land Use Research Institute, 1995.

van Beek, P. and J. Nunn. *Introduction to Rapid Appraisals*. SyTREC Pty Ltd, 1995.

van Veldhuizen, D. A. and G. B. Lamont. Multi-objective evolutionary algorithms: Analysing the state of the art. *Evolutionary Computation*, 8(2):125–147, 2000.

Zitzler, E. and L. Thiele. An evolutionary algorithm for multiobjective optimisation: the strength pareto approach. TIK-Report 43, Institute für Technische Informatik and Kommunikationsnetze, Zürich, 1998.