

Exploring the Geographic Consequences of Public Policies Using Evolutionary Algorithms

David A. Bennett,* Ningchuan Xiao,** and Marc P. Armstrong*

*Department of Geography, The University of Iowa

**Department of Geography, The Ohio State University

Public policies with geographical consequences are often difficult to analyze because they affect multiple stakeholders with competing objectives. While such problems fall conceptually into the domain of multiobjective evaluation, associated analytical techniques often search for a single optimum solution. Within the context of geographical problems, optimality often means different things to different stakeholders and, thus, an *optimum optimum* may not exist. In this article, we present a new technique based on an evolutionary algorithm (EA) that produces a large number of optimal and near-optimal solutions to a large class of land management problems. As implemented for this article, solutions represent landscape patterns that produce services that meet stakeholder needs to varying degrees. The construction of curves that illustrate the trade-offs among various services given limited resources is central to this approach. Decision makers can use these curves to help find solutions that strike a balance among conflicting objectives and, thus, meet stakeholder needs. To provide context to this work we consider the impact of the U.S. Department of Agriculture's (USDA) Conservation Reserve Program on rural landscapes. Three objectives are assumed: (1) maximize farm income, (2) maximize environmental quality, (3) minimize public investment in conservation programs; the first two are viewed as services desired by stakeholders. Analytical and visualization tools are developed to reduce the burden associated with exploring the large number of solutions that are produced by this technique. The results illustrate that the EA-based approach can produce results equal to and significantly more diverse than conventional integer programming techniques.

Key Words: spatial evolutionary algorithms, multiobjective optimization, decision support, agricultural policy.

Introduction

In the U. S., where private property rights are paramount, policymakers often must rely on the voluntary cooperation of a large number of independent land managers to achieve public goals. Public policies must be so crafted as to entice land managers into behaving in a particular way while they simultaneously strive to meet their own objectives. Agricultural policies, for example, influence, but typically do not dictate, farmers' decisions about where to farm, what to grow, and how to manage their land. Whether driven by vocation or avocation, different groups of land managers (stakeholders) are likely to disagree on the relative importance of public and private objectives, and policymakers must create policies that meet the needs of this diverse clientele.

To facilitate the evaluation of public policies we frame them within three interrelated analytical spaces: (1) solution, (2) geographic, and (3) objective (Figure 1). The solution space is comprised of all feasible policy or regulatory scenarios, each of which has intended, and often unintended, consequences. These consequences are often made manifest in geographical space. A new policy, for example, could precipitate changes in land

cover and management or influence spatial patterns of access and transport. Geographical and nongeographical consequences alike can often be transformed into indices that gauge how well particular policies meet stated objectives. The collective set of indices produced by the solution space constitutes the objective space. Collectively, we refer to this set of three linked spaces as the decision space of public policies with geographical consequences. The practice of searching for an optimal policy, however, is often complicated because:

1. the objective space is multimodal,
2. geographical problems are combinatorially complex,
3. decision makers represent competing stakeholder groups with distinctly different objectives,
4. all objectives may not be quantifiable,
5. all stakeholders may not act in an optimal manner, and
6. alternatives that are similar in one space can be dissimilar in one (or both) of the other two spaces.

The search for an optimal public policy can, therefore, be classed as a "wicked," or semistructured, problem as not all objectives can be represented in mathematical form

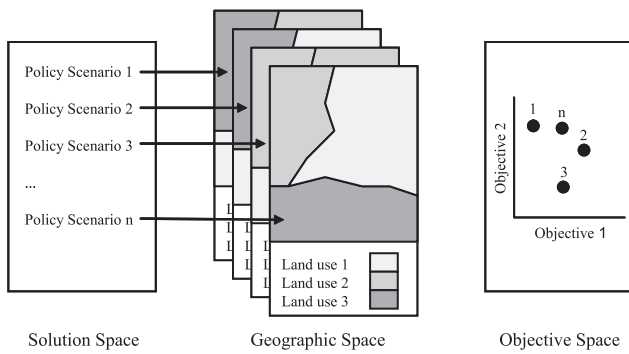


Figure 1. The decision space for public policies with geographic consequences is comprised of the solution space, geographic space, and objective space.

(Rittel and Weber 1973; Sprague and Carlson 1982; Densham 1991). Furthermore, it is likely that stakeholders will disagree on the relative importance of these objectives, and thus, optimality means different things to

different people. While an *optimum optimum* may not exist for such problems, policymakers must, nevertheless, make decisions.

In this article we present a new methodology that is designed to support the analysis of public policies with geographical consequences. Our method employs an evolutionary algorithm (described below) that operates on a digital representation of geographic space. We compare and ultimately integrate this approach with a more conventional method based on integer programming (Figure 2). This integration is accomplished by using the output of the integer programming method to seed the evolutionary algorithm. To motivate the presentation and discussion of this new method we place our work into the context of U.S. agricultural policy. In this context, we view the landscape as an engine of production capable of providing ecosystem services that meet stakeholder objectives. These services include, but are not limited to, agricultural income and environ-

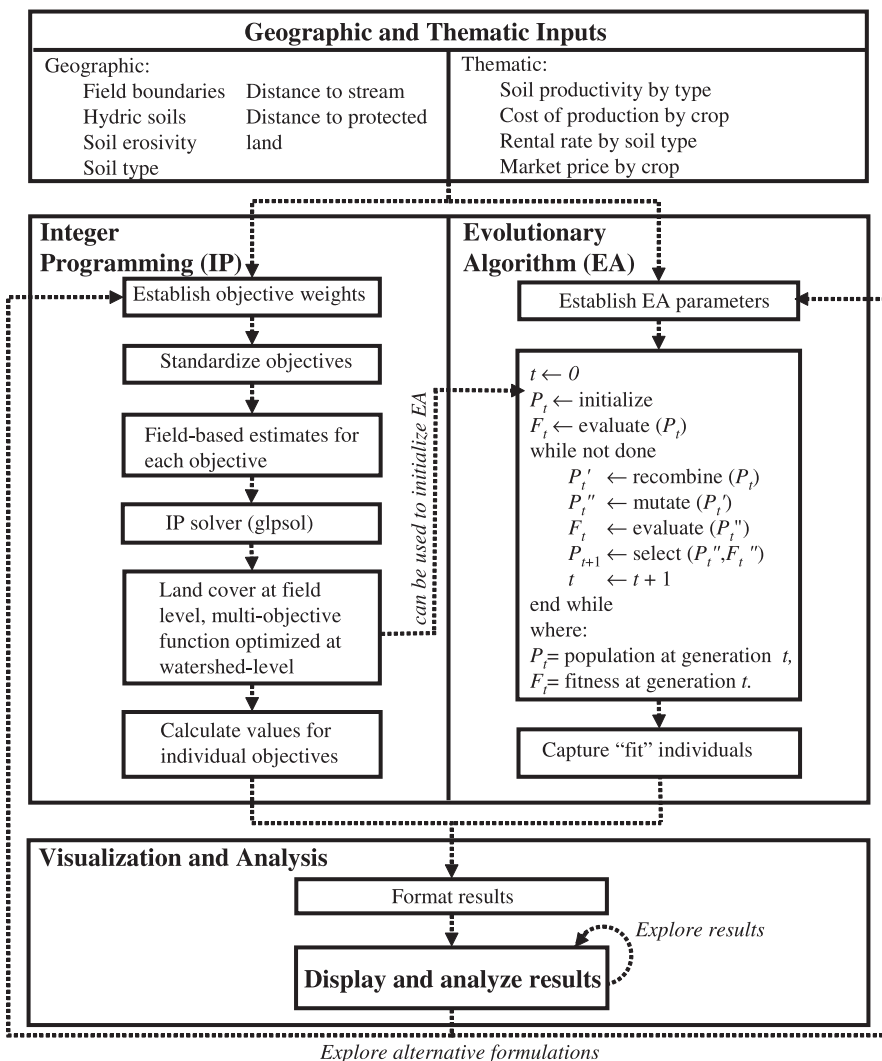


Figure 2. Basic elements of the presented approach.

mental quality. Public and private decision makers must decide how land and financial resources should be allocated among such services. Our goal is to support this decision-making process by producing: (1) alternative solutions that lie along a production possibility frontier; (2) visualization techniques that link the solution, objective, and geographic spaces; and (3) tools that help policy analysts identify interesting solutions.

Evaluating Multiobjective Geographical Problems

We begin with the assumption that most public policies are designed to meet multiple objectives and, as such, multiobjective evaluation techniques provide a logical methodological framework for policy evaluation. There are, however, several complicating factors associated with traditional methods of multiobjective evaluation (Cohon 1978; Miettinen 1999; Mimouni, Zekri, and Flichman 2000). First, a technique must be found to integrate incommensurable criteria. Multiobjective evaluation techniques often collapse multiple objectives into a single objective function to perform this integration, a process known as scalarization. To accomplish this task, decision makers must know, and agree upon, the relative importance of competing objectives (Malczewski 1999; Coello 2000). Though several methods have been developed to assist in the process of scalarization (e.g., see Saaty 1980; Carver 1991; Jankowski 1995; Malczewski 1999), all of them have significant limitations. Reaching a consensus on the relative importance of alternative criteria is, for example, difficult when decision makers possess distinctly different views of a problem. Furthermore, scalarization requires that the values associated with specified objectives be standardized to ensure that the analysis remains unbiased by objective-specific units (e.g., dollars versus parts per million). Standardization also requires a priori knowledge of the maximum and minimum obtainable objective values, and this knowledge is sometimes difficult to obtain.

Once scalarized, traditional optimization techniques can be used to find an “optimal” solution. These techniques, however, have several limitations. First, they are often based on assumptions that may not hold true in the context of geographical problems (e.g., a convex objective space). Second, the computational complexity of related algorithms often renders real-world problems intractable. Third, these approaches converge toward a single solution (which can be a local optimum). Techniques designed to converge toward a single solution, however, may fail to identify interesting alternatives

because, as we demonstrate in this article, positions close together in objective space often have significantly different geographic manifestations. Furthermore, the semi-structured nature of many spatial problems suggests the need for techniques that produce a wide array of potential solutions, as well as for tools that help decision makers evaluate trade-offs among conflicting objectives (Brill et al. 1990). Evolutionary algorithms can be adapted to produce a diverse set of alternative solutions that represent distinct trade-offs among objectives.

Evolutionary Algorithms

During the past decade, evolutionary algorithms (EA) have gained widespread recognition for their role in multiobjective optimization (see Deb 2001). EAs are based on an evolutionary metaphor. At the start, a population of solutions is created, and through selective pressure and the manipulation of digital genetic material, this population evolves over successive generations toward optimal solutions. The genetic material of individuals in this population is defined by a set of distinguishing characteristics (their genotype), which is often implemented as a one-dimensional array of values that acts as a chromosome. In the context of multiobjective optimization, the set of objectives used by decision makers to evaluate alternative solutions provides the selective pressure. A solution’s evolutionary *fitness* is measured by how well it meets these objectives (the response of a phenotype to selective pressures) (Bäck 1996; Bäck, Hammel, and Schwefel 1997). Several researchers have found that EAs are robust and can successfully evolve optimal and near-optimal solutions to multiobjective problems (Deb 2001; Xiao, Bennett, and Armstrong 2002; Armstrong, Xiao, and Bennett 2003). They are, however, heuristic devices used to explore large solution spaces and, thus, can become trapped in local optima due to premature convergence. The results of EAs are, therefore, often referred to as the “best solution(s) found so far.”

The genetic make-up that defines the initial population is often created through a random process, although, as we will demonstrate, prior knowledge about the problem and heuristics can improve the performance of the algorithm. New solutions are produced by an EA through the application of genetic operators; the two most common operators are recombination and mutation. Recombination incorporates characteristics of two parent solutions into one or more progeny. Individuals with high fitness values typically have a higher probability of being selected for recombination and, thus, a higher probability of contributing to the next generation of solutions.

Recombination, therefore, is an attempt to exploit successful adaptations found in the known solution space. Mutation operators, on the other hand, randomly modify part of the genetic material of a single individual. These operators are implemented to force the search process into unexplored regions of the solution space. In some implementations, a fixed percentage of fit individuals is copied into the next generation without modification. This procedure, referred to as *elitism*, ensures that the best solutions found so far are not lost from the population.

Representational Form and Evolutionary Algorithms

EA is a generic term for a family of four archetypes, derived from the basic evolutionary metaphor outlined above, that are distinguished by representational form, objective, and evolutionary strategy (e.g., an emphasis of recombination versus mutation). These archetypal forms are: genetic algorithms (Holland 1975; Goldberg 1989), evolutionary strategies (Rechenberg 1965), evolutionary programming (Fogel 1962), and genetic programming (Koza 1992). Genetic algorithms traditionally limit the representation of the chromosome to a binary vector and rely most heavily on recombination strategies to evolve better solutions; it is an exploitive approach. Evolutionary strategies (ES) and evolutionary programming (EP) explicitly support integer and floating-point representations. Adaptation in ES and EP is driven by mutation; these are explorative approaches. Finally, genetic programming techniques are used to produce sets of rules or statements (e.g., a computer program) that generate desired outcomes. In practice, it is often necessary to fit the form of an EA to the problem being addressed, and representations often borrow concepts from more than one of the above archetypal forms.

Evolutionary Algorithms and Spatial Analysis

EA methodologies have been applied successfully to a variety of geographical problems. Chambers and Taylor (1996), Bennett, Wade, and Armstrong (1999), Matthews, Sibbald, and Craw (1999), and Xiao, Bennett, and Armstrong (2000), for example, have illustrated the utility of EAs in the context of environmental analysis. Hosage and Goodchild (1986), Dibble and Densham (1993), Krzanowski and Raper (1999), Brookes (2001), and Xiao, Bennett, and Armstrong (2002) have applied EAs to locational problems and Balling et al. (1999) used them to identify urban development patterns that minimized traffic congestion. Even in the context of geography, however, evolutionary computation is often used only as a heuristic to evaluate well-defined problems.

Less attention has been given to the representation of semistructured problems and the elicitation of new spatial solutions to such problems, although we suggest that it is in this role that EAs will prove most beneficial in geographical analysis.

The work presented in this article is an extension of research described by Bennett, Wade, and Armstrong (1999), which applied evolutionary computation to elicit compromise solutions for a land resource allocation problem. Their approach used the Morton indexing scheme to "linearize" space (Samet 1990, 14) into a chromosome that could be manipulated by traditional EA operators. In particular, a two-point recombination operator exchanged land cover within a randomly selected region of two parent landscapes. Similarly, a mutation operator placed a new land cover into a region of a selected landscape (see Bennett, Wade, and Armstrong 1999 for details). The evolving population in this approach was comprised of 100 different landscapes, and the objective was to maximize a scalarized multiobjective evaluation function. The specified objectives were to minimize soil erosion, maximize wildlife habitat potential, and maximize profit. Three software agents were created to represent the interests of three stakeholder groups (the farming community, conservation groups, and wildlife enthusiasts), and the EA was used to search for compromise solutions that would be acceptable to these competing agents. A rank-based (mean agent-specific rank) selection algorithm was implemented, and the solution with the highest fitness value was propagated into the next generation (i.e., *elitism* was used).

Because the relative spatial position of land cover types is an important determinant of habitat quality, the problem described by Bennett, Wade, and Armstrong (1999) was computationally intractable using brute-force methods. Their methodology, however, quickly evolved land-use patterns that were logical at the landscape scale and, thus, their results suggested that EAs could help decision makers explore difficult geographical problems. That experiment also helped to elucidate several avenues for future research that are addressed in this article. First, the Morton indexing scheme used to represent spatial pattern failed to identify and exploit fine-grained features in the landscape. For example, it had difficulty finding narrow, meandering stream corridors, which can have considerable environmental significance. Second, the resulting population was comprised of relatively homogeneous landscapes, though maintaining a variety of distinctly different alternatives that meet stated criteria is often deemed to be more useful during decision-making processes (Brill et al. 1990). Third, the evolutionary algorithm was designed

specifically to evolve toward a single point in the objective space. This was appropriate given the problem presented by Bennett, Wade, and Armstrong (1999) because their goal was to find compromise solutions in a complex problem domain. In the broader context of decision support, however, it may be more desirable to search widely throughout the decision space and to investigate the trade-offs that can be made as finite resources are used to produce desired outcomes.

The Production Possibility Frontier

A production possibility frontier (PPF) illustrates the combinations of outputs that can be produced given a finite set of resources, and is comprised of the set of nondominated solutions. Given two solutions (x and y) x dominates solution y if and only if:

- x is at least as good as y for all objectives, and
- x is strictly better than y for at least one objective.

More formally, for a problem of $\max(f_1, f_2, \dots, f_k)$ x dominates $y \iff \forall i, f_i(x) \geq f_i(y) \wedge \exists j, f_j(x) > f_j(y)$ where:

- f_i = objective function for criterion i ,
- $i, j \in \{1, 2, \dots, k\}$, and
- k = total number of criteria.

In the problem investigated here, the resources are land and tax dollars, which can be used to produce two forms of ecosystem service, food crops (and thus agricultural income), and environmental quality. The work of Armstrong, Xiao, and Bennett (2003), Deb (2000), Fonseca and Fleming (1993, 1995), and Goldberg (1989) illustrates how EA strategies (typically GA/EP hybrids) can be used to search widely throughout a solution space and identify those solutions that lie at or near a PPF. Figure 3 illustrates the difference between an EA approach that collapses many objectives into a single objective function (scalarization) and one designed to produce a PPF. While the population may start out diversified in both approaches, the goal of scalarization is to evolve toward a single point along the frontier. This point represents a particular trade-off between two or more services. If decision makers know a priori the balance that they wish to strike among services, then this can be a useful approach. In contrast, the goal of constructing a PPF is to produce a diversified population of solutions that will illustrate: (1) trade-offs among competing services and (2) the domain of feasible and infeasible solutions. Note that in this approach, unlike many optimization procedures, dominated solutions near the front are deemed useful and are purposefully maintained because:

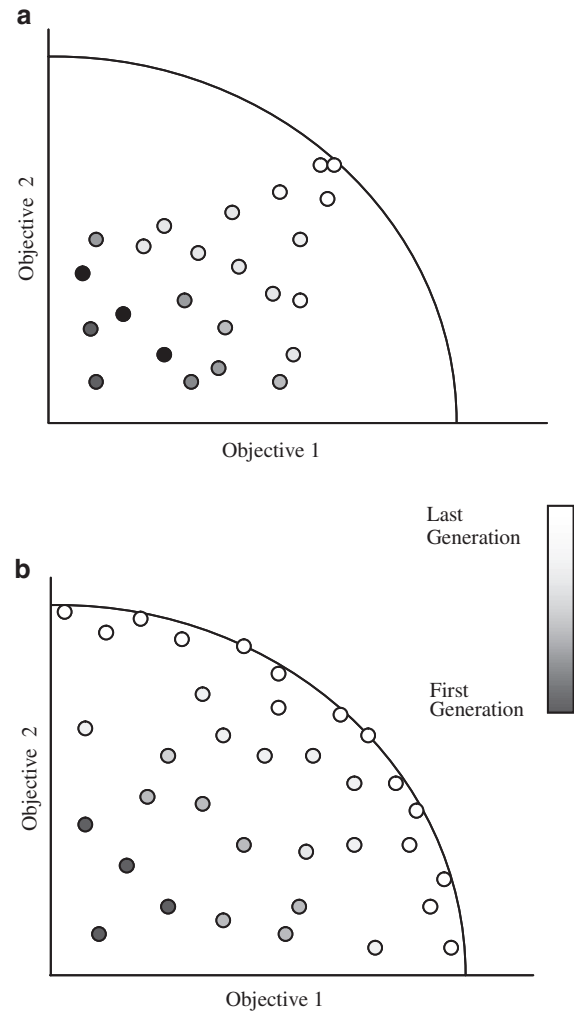


Figure 3. 3a represents the evolutionary trajectory for scalarized multiobjective fitness functions, 3b the trajectory for a production possibility frontier.

1. they provide “genetic” diversity that may prove useful as the evolutionary process proceeds, and
2. the semistructured nature of many geographical problems precludes a deterministic solution to “optimality.”

Study Context

We focus on the Big Creek watershed in the southernmost region of Illinois to illustrate the utility of our approach (Figure 4). The watershed covers approximately 13,400 hectares and is associated with two distinct physiographic features. The upper part of the watershed is in the Central Lowland. This part of the watershed is hilly and the land cover consists of grains intermixed with forest and grassland. The lower part of

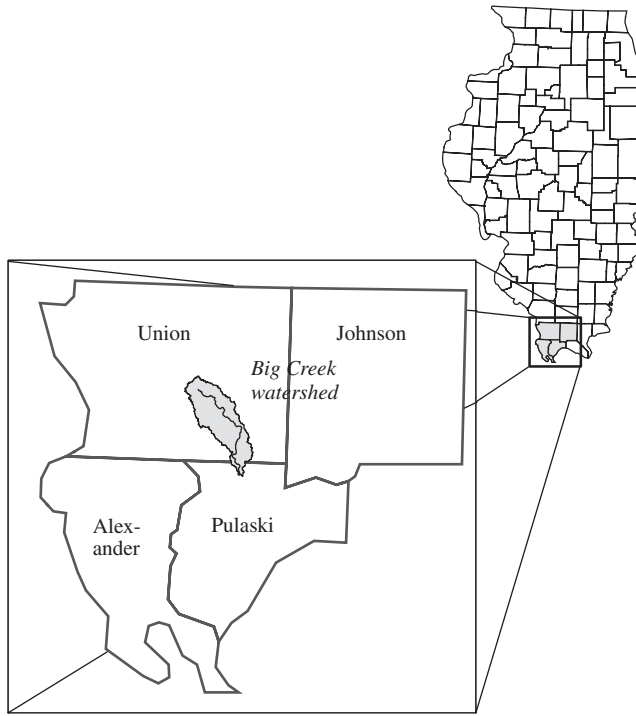


Figure 4. This study was conducted in the Big Creek watershed in southern Illinois.

the watershed opens into the broad, flat Cache River watershed. The Cache River lies along an abandoned section of the Ohio River and is at the northern-most edge of the Mississippi Alluvial Plain. This lower area is farmed intensively. Big Creek empties into the Cache River just downstream of Buttonland Swamp, which is the centerpiece of a conservation effort being undertaken by the U.S. Fish and Wildlife Service, U.S. Forest Service, the Illinois Department of Conservation, and the Nature Conservancy. Buttonland Swamp serves an important ecological role within the region as a wintering ground for migratory waterfowl, a migratory corridor for neotropical songbirds, and as a haven for a variety of threatened and endangered species. During periods of peak discharge, water and entrained sediment flows back into Buttonland Swamp from Big Creek and, therefore, the management of the Big Creek watershed is of considerable interest to a variety of stakeholders with competing objectives (Bennett et al. 2001).

We consider here the impact of the conservation reserve program (CRP) on this watershed. The CRP pays farmers to remove farm fields from agricultural production and is administered by the U. S. Department of Agriculture (USDA). Farmers interested in enrolling a field into the CRP must enter into a bidding process. The USDA selects a subset of all fields offered for enrollment by farmers based on cost and perceived environmental

return. Embedded within the CRP legislation (P.L. 99–198) are implementation-specific details regulating, for example, the maximum acceptable area of land that can be enrolled in a county (25 percent). However, a broader set of regulatory rules could be used to define the CRP. For example, decision makers may wish to establish policies that define acceptable thresholds for:

1. economic impact,
2. public investment, or
3. environmental return.

It is at this level of policymaking that we focus our analysis.

Methodology

To illustrate our approach we explore the trade-offs among three policy objectives: (1) maximize the gross marginal return (total revenue minus variable costs) of agricultural production; (2) maximize the environmental benefit derived from public investment in the CRP; and (3) minimize public investment in the conservation reserve program. Payments made to farmers as part of the CRP are not considered when calculating marginal return; to do otherwise would underrepresent the trade-off between environmental quality and agricultural production. Given this formulation of the problem, we expect that as public investment in the CRP increases, environmental benefit will increase and the gross marginal return associated with agricultural production will decline as land is taken out of production and placed under various forms of conservation management. We explore the decision space associated with the CRP using a standard integer programming (IP) approach and an EA. As specified by current regulations, we also restrict CRP enrollment to 25 percent of the landscape.

Maximize Gross Marginal Return

Gross marginal return for the Big Creek Watershed is calculated as:

$$O_1 = \sum_j \sum_k x_{jk} a_j (f_{jk} p_k - c_k), \quad (1)$$

where:

- O_1 = objective 1,
- j = index of farm fields ($0 \leq j \leq 961$),
- k = index of crop cover types ($0 \leq k \leq 4$),
- $x_{jk} = \begin{cases} 1, & \text{if field } j \text{ has cover type } k \\ 0, & \text{otherwise,} \end{cases}$
- a_j = area of field j ,

f_{jk} = productivity of field j for cover type k (areally weighted average of all soil types within field j),
 p_k = price for cover type k , and
 c_k = variable costs of production per unit area for cover type k .

Agricultural commodity prices (p_k) were set at the 1993 to 1997 five-year mean (Farmdoc 2003). These dates were used so that model results can be compared to cropping practices documented in the 1997 Census of Agriculture. Data for the variable costs of crop production in southern Illinois are available from 1998 to 2003; these values, however, remained relatively constant. The values used for prices and costs are shown in Table 1.

Table 1. Cost and Revenue for Modeled Crop Types (Farmdoc 2003)

Cover ID	Crop	Cost/ha	Price
0	Corn	\$400.31	\$2.74/bu
1	Soybean	\$259.46	\$6.65/bu
2	Wheat	\$229.81	\$3.59/bu
3	Hay	\$301.47	\$68/ton
4	Double crop	Function of soybean and wheat	

Maximize Environmental Benefit

Environmental benefit is estimated using USDA guidelines for calculating the environmental benefit index (EBI) for Union County, Illinois (see Ribaud et al. 2001 for a detailed discussion on the use of the EBI in the CRP). The EBI is used to guide the CRP enrollment process and is comprised of seven major elements that are summed to calculate an overall index:

$$EBI = N1 + N2 + N3 + N4 + N5 + N6 + N7, \quad (2)$$

where:

- N1 = benefit to wildlife habitat (0–100 points),
- N2 = benefit to water quality attributable to reduced erosion, runoff, and leaching (0–100 points),
- N3 = on-farm benefit attributable to reduced erosion (0–100 points),
- N4 = enduring benefits (0–50 points),
- N5 = air quality benefits from reduced wind erosion (0–35 points),
- N6 = benefits of enrollment in conservation priority areas (0–25 points), and
- N7 = cost of implementing CRP management activity (point value not established by regulation).

Many of these elements are decomposed into more specific criteria. For example, the benefit to wildlife is calculated as:

$$N1 = (N1a/50) * (N1a + N1b + N1c + N1d + N1e + N1f), \quad (3)$$

where:

- N1a = cover (0 to 50 points),
- N1b = endangered species (0 to 15 points),
- N1c = proximity to permanent water (0, 5, or 10 points),
- N1d = adjacency to protected areas (0, 5, or 10 points),
- N1e = wildlife enhancements—food plots, wetland restoration (0 or 5 points), and
- N1f = restored wetland and upland cover (0 or 10 points).

In this research, EBI values were not estimated for N5 and N7 (N5 is not relevant to this study area and data for N7 were not available). Figure 5 illustrates the pattern of four spatial variables that are included in this analysis.

Eleven different CRP practices are modeled, and each practice is associated with an EBI value (Table 2). The watershed level EBI value is calculated as:

$$O_2 = \sum_j \sum_l a_j x_{jl} ebi_l \quad (4)$$

where:

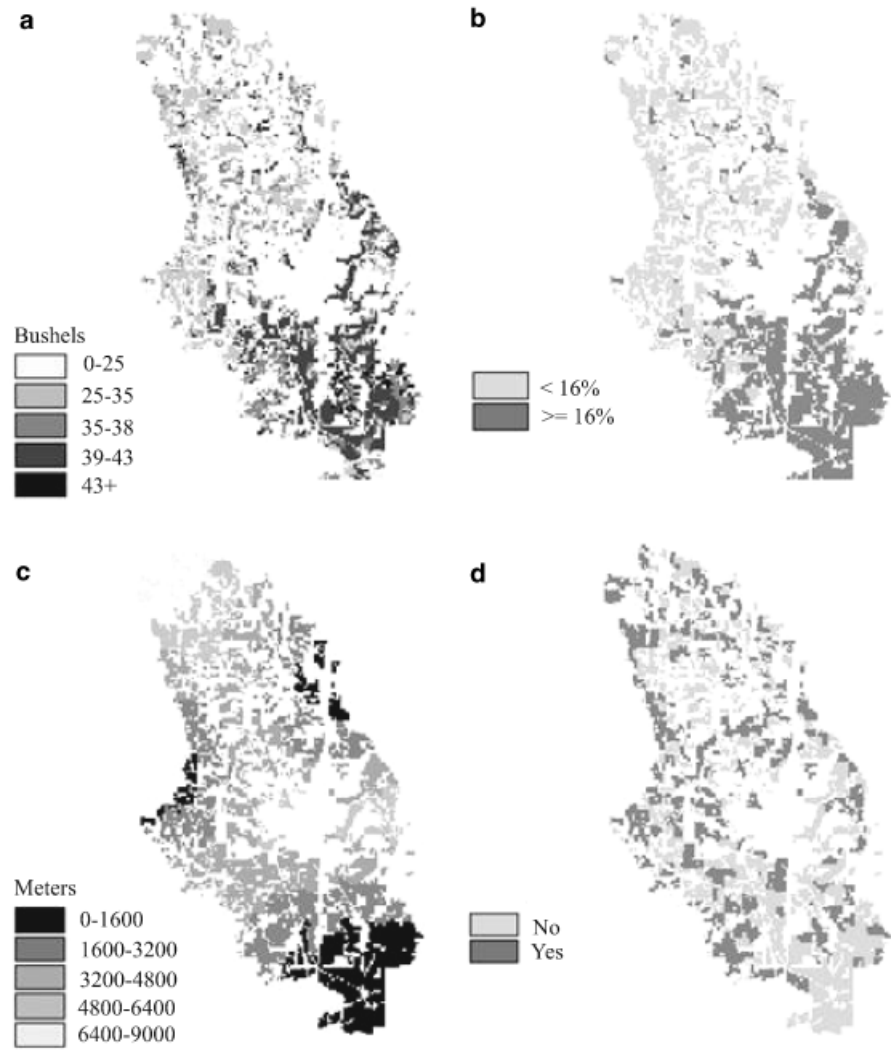
- O_2 = objective 2,
- j = index of farm fields ($0 \leq j \leq 961$),
- l = index of CRP cover types ($5 \leq l \leq 15$),
- a_j = area of field j ,
- $x_{jl} = \begin{cases} 1, & \text{if field } j \text{ has CRP cover type } l \\ 0, & \text{otherwise, and} \end{cases}$
- ebi_l = EBI value for CRP cover type l .

Minimize Public Investment

The CRP rental rate is the cash payment made per unit area to farmers for removing land from production. This rate is a function of soil type, and for each farm field it is calculated as the areally weighted average of the rates associated with all soil types within that field. CRP rental rates were provided by USDA district conservationists from the Illinois counties of Union, Johnson, and Pulaski. Public investment is, therefore, calculated as:

$$O_3 = \sum_j \sum_l a_j x_{jl} inv_j \quad (5)$$

Figure 5. The spatial pattern of four important inputs: (a) soil productivity for crops, input to marginal return; (b) hydric soils, input to EBI parameters N1 and N2; (c) distance to protected land, input to EBI parameter N1, (d) erodibility, input to EBI parameters N2 and N3.



where:

O_3 = objective 3,

j = index of farm fields ($0 \leq j \leq 961$),

l = index of CRP cover types ($5 \leq l \leq 15$),

a_j = area of field j ,

$x_{jl} = \begin{cases} 1, & \text{if field } j \text{ has CRP cover type } l \\ 0, & \text{otherwise, and} \end{cases}$

inv_j = Public investment needed to enroll farm field j .

Integer Programming (IP)

An IP model was developed to produce a set of baseline solutions that were used to test the efficacy of the EA developed as part of this research. The three objectives listed above are integrated into a single scalarized multiobjective function, which is maximized using the *glpsol* software in the GNU linear programming

kit (*glpk*). More formally, we attempt to:

$$\text{Maximize: } \sum_m w_m O_m / s_m \quad (6)$$

Subject to:

$$\sum_j \sum_n x_{jn} = 1 \quad (7)$$

$$\sum_j \sum_l a_j x_{jl} \leq 0.25 \sum_j a_j \quad (8)$$

where:

j = index of farm fields ($0 \leq j \leq 961$),

m = index of objectives ($1 \leq m \leq 3$),

n = index of all cover types ($0 \leq n \leq 15$),

l = index of CRP cover types ($5 \leq l \leq 15$),

w_m = weight of objective m ,

O_m = the m^{th} objective (Equations 1, 4, and 5),

s_m = normalization for the m^{th} value,

Table 2. CRP Practices

Cover ID	Description	EBI Pts.
5	Introduced grass	10
6	Mixed stand of 3 or more plant species, with at least 1 introduced grass, and 1 forb or legume species	30
7	Mixed stand of 4 or more plant species, with at least 2 introduced grasses and 1 forb or legume	40
8	Solid stand of 1–3 native species	20
9	Mixed stand of 4 or more plant species, with at least 2 native grasses and 1 shrub, forb, or legume	40
10	Mixed stand of 5 or more plant species, with at least 2 native grasses and 1 shrub, forb, or legume	50
11	Solid stand of nonmast-producing hardwood species	20
12	Solid stand of a single hard mast-producing species	40
13	Mixed stand of hardwood species best suited for wildlife	50
14	Two shrub or trees species and at least 2 native	40
15	Wetland restoration	50

$$x_{jn} = \begin{cases} 1, & \text{if field } j \text{ has cover type } n \\ 0, & \text{otherwise, and} \end{cases}$$

$$a_j = \text{area of field } j.$$

The first constraint (Equation 7) ensures that one and only one cover type is associated with each field, the second constraint (Equation 8) ensures that no more than 25 percent of the landscape is placed into the CRP as required by USDA guidelines. The value of s_m is set equal to the maximum value obtained from the EA and single-objective IP runs.

Using this approach, we attempted to locate a number of points that would define the PPF by manipulating the objective weights associated with the three stated objectives (Table 3). The number of points required to accomplish this task is difficult to estimate without a priori knowledge of the PPF's form. Seven points were selected to explore the shape of the front. Since neither the IP nor EA results suggested that discontinuities or

concavities existed in the PPF, the analysis of additional points was not deemed necessary. Other problems, however, might require a larger number of points to define a PPF.

Evolutionary Algorithm

With this research we extend the work of Bennett, Wade, and Armstrong (1999) and Armstrong, Xiao, and Bennett (2003) to produce landscapes that fall along a PPF that is defined by the trade-offs among competing objectives. The decision to place a particular tract of land into the CRP or to plant crops is influenced by soil productivity, commodity prices, the CRP rental rates, the costs of production, and expected environmental return. The chromosome operated on by the EA is represented as an integer array. The index value associated with each element in this array links it to a specific farm field, and the value stored at this index location refers to the associated land cover type (Figure 6). The decision variable used to optimize landscape level indices, therefore, is land cover at the field level. Conceptually, this approach is similar to that of Balling et al. (1999). In each approach the chromosome contains pointers to geographic regions (here agricultural fields, land use zones in Balling et al. 1999). Furthermore, both techniques are designed to highlight trade-offs among competing objectives. How the EA is implemented, however, differs significantly between these two approaches. For example, the Balling et al. (1999) approach does not include techniques to: (1) promote diversification along the PPF, (2) validate results, or (3) visualize the decision space.

The Fitness Function. The fitness function must move the solution process toward the PPF and promote a diverse solution set to ensure that all portions of the curve are fully developed (Deb 2001). Goldberg (1989) developed a fitness evaluation method that assigns all nondominated individuals in the current generation the rank value of 1 (Figure 7). These individuals are removed from the evaluation and a new set of nondominated individuals is identified, which are assigned a rank value of 2. This process continues until all individuals are ranked or until a user-defined threshold is reached (e.g., ignore all individuals below a rank of 4). Once each individual is ranked, fitness is calculated as a function of $1/\text{rank}$. While such a criterion effectively moves the population toward the front, it does not explicitly ensure diversity, and the algorithm can converge to subregions along the front.

Table 3. Objective Weights for Integer Programming Runs

Run	Gross Margin	Environmental Benefit	Public Investment
1	0.8	0.1	0.1
2	0.6	0.2	0.2
3	0.4	0.3	0.3
4	0.334	0.333	0.333
5	0.3	0.4	0.3
6	0.2	0.6	0.2
7	0.1	0.8	0.1

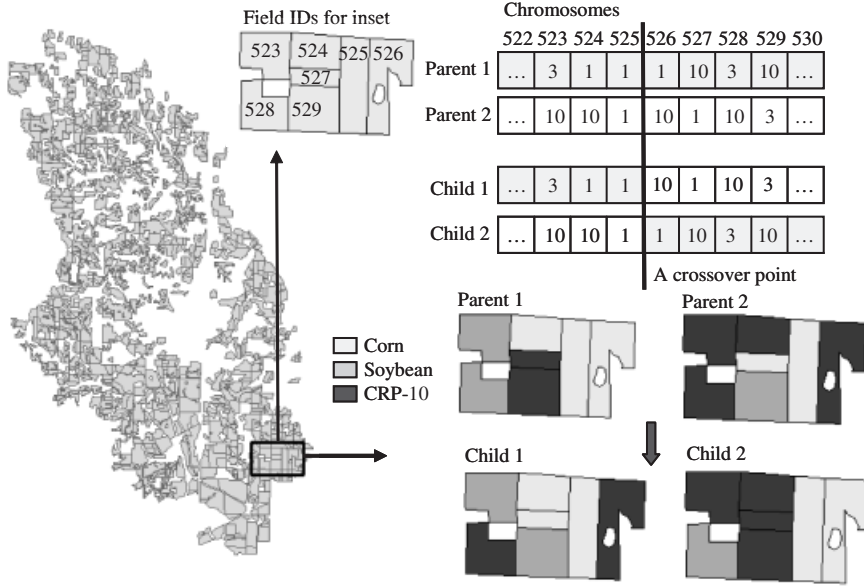


Figure 6. The recombination operator swaps the land cover associated with a random number of contiguous fields between two parent landscapes to produce two new solutions.

To overcome this problem we have implemented two techniques that are designed to promote solution diversity in EAs. First, we implemented Srinivas and Deb's (1994) nondominated sorting algorithm. Central to this algorithm is the concept of a *niche count*. The niche count is designed to boost the fitness of those solutions that lie along sparsely populated regions of the emerging front:

$$nc_p = \sum_q Sh(d_{pq}), \quad (9)$$

$$Sh(d_{pq}) = \begin{cases} 1 - (d_{pq}/\sigma_{share})^\alpha, & \text{if } d_{pq} < \sigma_{share} \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

where:

nc_p = niche count for landscape p ,
 $Sh(d)$ = sharing function,




Figure 7. Goldberg's rank-based fitness algorithm is used to derive fitness values.

d_{pq} = normalized distance in objective space between landscapes p and q in the population,
 σ_{share} = neighborhood parameter, and
 α = exponent capturing the effect of distance on the share function.

As implemented here, distance is defined as the number of fields in landscape p that have a different cover type than the same field in landscape q . Past research suggests that the nondominated sorting algorithm is particularly sensitive to σ_{share} , the parameter that determines the neighborhood within which fitness is apportioned (Deb 2001). In theory, a smaller neighborhood leads to more diversified solutions and, thus, a more complete PPF. For the runs presented in the Results section, $\sigma_{share} = 0.15$ and $\alpha = 2$. The sensitivity of the EA to these variables is investigated below.

The nondominated sorting algorithm is implemented by calculating, in sequence (adapted from Deb 2001, 2003):

1. values for O_1 , O_2 , O_3 , as defined above for all solutions q in population P ,
2. initial fitness value F'_q for all solutions q in population P by implementing Goldberg's nondominated sorting algorithm,
3. niche count, nc_q , for all solutions q in population P , and
4. shared fitness for all solutions q in population P as $F''_q = F'_q / nc_q$.

Our second approach creates several subpopulations (Bennett, Armstrong, and Wade 1996; Armstrong, Xiao, and Bennett 2003; Xiao and Armstrong 2003), with the

goal of evolving specialized subspecies of solutions that are well adapted to overlapping segments of the PPF. To accomplish this process of “allopatric speciation” each subpopulation is subjected to somewhat different selective pressures and the exchange of genetic material among subpopulations is regulated by user-specified migration rules (i.e., rules that determine the subset of subpopulations to which an individual can “move”). For the results presented below, the following six subpopulations are defined:

- Subpopulation 1 weights all objectives evenly,
- Subpopulation 2 attempts to maximize farm income and EBI,
- Subpopulation 3 attempts to maximize farm income and minimize public investment,
- Subpopulation 4 attempts to maximize EBI and minimize public investment,
- Subpopulation 5 attempts to maximize farm income, and
- Subpopulation 6 attempts to maximize EBI.

The objectives associated with these subpopulations (e.g., maximize farm income and EBI) were defined in this manner to encourage the evolution of a diverse population that extends across the entire PPF. Repeated runs of the system with various settings suggested that a population of 100 individuals that was allowed to evolve over 500 generations produced consistent results. Consequently, each subpopulation was instantiated with those values.

Initialization. We experimented with three different techniques for initializing individuals in the population: modified random, heuristic, and seeding the EA with IP results. The initial population developed by the modified random initialization procedure contains landscapes with:

1. 20 percent probability that all fields have the same randomly selected crop cover type,
2. 20 percent probability that all fields have the same randomly selected CRP cover type, and
3. 60 percent probability that each field will contain a randomly selected cover type.

Producing landscapes comprised of the same randomly selected cover type (1 and 2 above) provides a reasonable expectation that all cover types will be represented in all fields, thus providing needed genetic diversity.

The above algorithm follows the EA tradition of random initialization, but fails to exploit known properties of this particular problem. In an effort to produce a

more effective initial population (defined by time to convergence and the degree to which the PPF is produced) we use field-level values for EBI and gross marginal return as heuristics for watershed-level performance. This modified population contains:

1. eight individuals where all fields are set to the crop type that returns the maximum gross marginal return,
2. eight individuals where all fields are set to the CRP type that returns the maximum EBI value,
3. thirty-four individuals where each field has an equal probability of receiving the cover type that returns the maximum value for gross marginal return or EBI, and
4. fifty individuals produced in the manner outlined above for modified random initialization.

While the number of individuals in each of the four classes is somewhat arbitrary, the inclusion of multiple versions of the same individual increases the probability that useful genes will survive to subsequent generations.

We would expect even greater efficacy if the IP runs discussed above are used to initialize the population. Consequently, an initial population was produced using this EA/IP hybridized approach that contains:

1. five individuals produced by IP Run 1,
2. five individuals produced by IP Run 2,
3. five individuals produced by IP Run 3,
4. five individuals produced by IP Run 4, and
5. eighty individuals produced in the manner outlined above for random initialization.

Though the number of individuals in each of these five classes is, again, somewhat arbitrary, there was no evidence of an undeveloped PPF (e.g., premature convergence, gaps, or concavities). If the results suggested that the front was underdeveloped, then adjustments would have been made to this initialization procedure.

Genetic Operator. The recombination procedure implemented here selects two random integers from a uniform distribution ($U[0, \text{number of fields} - 1]$). These integers correspond to index values in the chromosome and, thus, farm fields in the study area. The recombination algorithm swaps the land cover of farm fields between these two indices (Figure 6).

The following three mutation operators were implemented to encourage exploration within the solution space:

1. randomly select n fields and change to a randomly selected cover type,

2. randomly select a block of contiguous fields and change to a randomly selected cover type, or
3. randomly select n fields and swap crop type to CRP land (or vice versa).

Each of these three mutation operators has an equal probability of being selected for execution given a user-defined mutation rate (set at 0.1 for the runs presented in the Results section).

Selection. A ranked-based, roulette-style selection algorithm (Goldberg 1989) was implemented to determine which individuals in generation t would be used to produce generation $t+1$. Twenty-five percent of the best individuals in generation t are copied into generation $t+1$ without modification, and the remaining 75 percent of the population (the gap size) is produced through recombination and mutation.

Visualizing the Decision Space

The PPF is designed to help decision makers understand the trade-offs associated with competing objectives. When the goal is to select a single solution on or near the frontier, however, the volume of alternative solutions associated with this approach can prove daunting (e.g., see Schwartz 2004). We have developed a set of tools to reduce the cognitive burden associated with the evaluation and selection of solutions. First, an interface was produced that links the three different spaces associated with the decision space (Figure 8, see also Xiao, Armstrong, and Bennett 2002). A table

contains information about the objective values associated with all landscapes, and it can be sorted by any objective. The sorting function allows analysts and decision makers to focus their attention on that subset of landscapes that lies within actual or proposed policy guidelines (i.e., the solution space can be constrained to meet user goals). This table is linked to a graphical view that presents the PPF (a representation of the objective space). By interactively selecting a point in the graph, the corresponding row in the table is highlighted (and vice versa). Once a landscape is identified in either the tabular or graphical view, a map can be produced (a representation of geographic space). We have also implemented two proximity indices that compare alternatives in the geographic and objective spaces. The Euclidean distance between two solutions in n -dimensional objective space is calculated as:

$$c_{pq} = \sqrt{\sum_m (v_{pm} - v_{qm})^2}, \quad (11)$$

where:

m = index of objective ($1 \leq m \leq 3$),

p, q = index of a landscape,

c_{pq} = distance between landscape p and landscape q in objective space, and

v_{pm} = value of landscape p for objective m .

The distance between two alternatives in geographic space is calculated as:

$$s_{pq} = \sum_j t_{ij}, \quad (12)$$

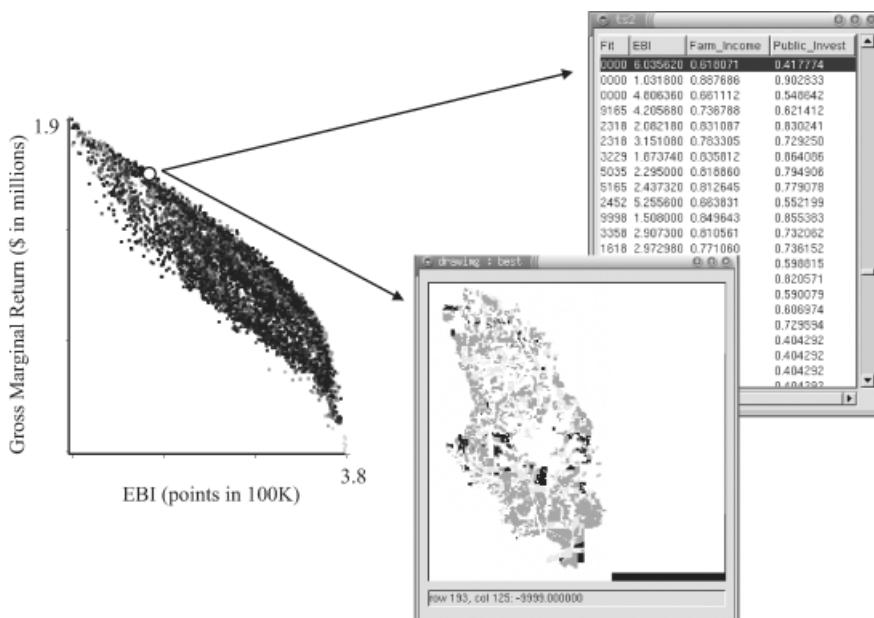


Figure 8. The user interface provides three views of the decision space, tabular, graphical, and geographical.

where:

$$s_{pq} = \text{distance between landscape } p \text{ and landscape } q \text{ in geographic space, and}$$

$$t_j = \begin{cases} 0, & \text{if land cover at field } j \text{ in landscape } p \\ & \text{is the same as land cover at field } j \text{ in} \\ & \text{landscape } q \\ 1, & \text{otherwise} \end{cases}$$

These two indices can be used to formulate queries such as:

1. select all solutions that are “near” to landscape p in the objective space, or
2. select all solutions that are “near” to landscape p in the geographic space,

where “near” is a user-defined proximity value. Furthermore, we can visualize “nearness to solution p ” in geographic space by symbolizing all q points along the PPF as a function of s_{pq} .

These indices can be used to drive additional queries of interest. For example, a user may want to find alternatives that are close to alternative p in the objective space, but distinctly different in the geographic space. To support such queries we first standardize s_{pq} and c_{pq} :

$$ss_{pq} = \frac{s_{pq} - s_{\min}}{s_{\max} - s_{\min}} \quad (13)$$

$$cc_{pq} = \frac{c_{pq} - c_{\min}}{c_{\max} - c_{\min}} \quad (14)$$

where:

- s_{\min} = the minimum s value between solution p and all other solutions,
- s_{\max} = the maximum s value between solution p and all other solutions,
- c_{\min} = the minimum c value between solution p and all other solutions, and

c_{\max} = the maximum c value between solution p and all other solutions.

We then calculate the following index:

$$S_q = \sum (w_{pq} ss_{pq} / cc_{pq}) / \sum w_{pq}, \quad (15)$$

where:

$$w_{pq} = \begin{cases} 0, & cc_{pq} > cc_o \\ 1, & \text{otherwise, and} \end{cases}$$

cc_o = user-specified proximity threshold in objective space.

Similarly, a user may want to find alternatives that are close to alternative p in the geographic space but distinctly different in the objective space. To support this class of query we calculate:

$$C_q = \sum (w_{pq} cc_{pq} / ss_{pq}) / \sum w_{pq}, \quad (16)$$

where:

$$w_{pq} = \begin{cases} 0, & ss_{pq} > ss_o \\ 1, & \text{otherwise, and} \end{cases}$$

ss_o = user-specified proximity threshold in geographic space.

Results

Table 4 documents the results of the integer programming base runs. Computation times for Runs 1 through 4 were quite fast, taking 6 seconds or less on a 2.4 Ghz Pentium IV computer with 512 MB RAM. As more weight is placed on environmental quality (Runs 5 through 7), however, the policy that limits total CRP land to less than 25 percent of the landscape becomes an important spatial constraint. In this situation the *glpsol* software failed to find a solution; the program terminated with an out-of-memory error after running for approximately three days. Though spatial constraints are commonly used

Table 4. Results of Integer Programming Runs

	1997 Census of Ag Union Co. IL	IP Run 1	IP Run 2	IP Run 3	IP Run 4
Crops as % of total		97.1%	96.8%	82.7%	75.1%
as % of total crops					
Corn	19%	11%	11%	11%	9%
Soybean	44%	62%	61%	58%	57%
Wheat	6%	4%	4%	5%	5%
Hay	31%	24%	24%	27%	28%
CRP as % of total		2.9%	3.2%	17.3%	24.9%
as % of total CRP					
10	NA	0%	0%	0%	25%
13	NA	52%	27%	40%	24%
15	NA	48%	73%	60%	51%

to create equity and manage costs in public policies, they often introduce considerable computational complexity, and the inability of IP algorithms to find a solution to such problems limits their utility in geographic analysis.

It is also worth noting that the landscapes produced by Runs 1 through 5 are comprised of CRP practices that return high EBI values. Practices rewarded with high EBI values, however, often require a high capital investment and produce a land cover that is difficult to revert to agricultural production (e.g., woodlands and wetlands). Experience shows that many farmers are reluctant to make such commitments and often opt for CRP practices that may have lower financial return but greater long-term flexibility. Furthermore, the IP approach produced landscapes that contain more soybean, and less corn than is suggested by the 1997 data. One reason for this situation is that the IP model did not consider the need to rotate soybeans with other crops in this region as a defense against soybean sudden death syndrome (Rupe, Robbins, and Gbur 1997). If corn and soybeans are aggregated into a single category, the percent distribution for all land cover categories is within 10 percent of the 1997 landscape. Given the uncertainties of agricultural production, we conclude that these landscapes provide a realistic benchmark against which the results of the EA approach can be compared.

Results From EA

Figure 9 illustrates the results of the evolutionary algorithm developed here. Each black dot in these graphics represents a landscape produced by the EA. Open circles are landscapes produced by IP Runs 1 through 4 and, thus, represent points along part of the PPF (note that IP Runs 5 through 7 failed to produce results and, thus, there are no open circles in the lower portions of the frontier). Each run of 500 generations required approximately 88 minutes to execute on a 2.4 Ghz Pentium IV computer with 512 MB RAM and produced 225,150 alternative landscapes. The 30 most-fit landscapes from each subpopulation in each generation are stored and presented in objective space (e.g., 45,000 landscapes are represented in Figures 9a, c, and e).

All three initialization algorithms produced fronts that are fully formed with solutions well diversified along the front. However, the population produced using the modified random initialization algorithm converged prematurely and, thus, failed to reach the PPF. In contrast, the populations produced by the heuristic and hybridized IP initialization algorithms produced fronts that reached the PPF (i.e., included the IP solutions). From the shading of the points (light representing earlier genera-

tions, dark representing later generations) we see that, generally speaking, the algorithm made steady progress toward the frontier. Note, however, that this progress is not necessarily monotonic. During the creation of each new generation, the processes of recombination and mutation can, and do, produce individuals that are less fit than those in previous generations. In fact, selecting a termination condition for EAs (e.g., stop after 500 generations) can prove challenging, and one indication that the evolution has run too long is that the population begins to devolve away from the front. In Figure 9e, for example, there is evidence that the hybrid IP algorithm operated too long. Note that in the lower portion of this graph light gray points (individuals from early generations) extend out beyond darker points (individuals from later generations). Figure 9f shows that the hybridized solution was, in fact, able to reach the front in 100 generations, while the heuristic initialization algorithm was not (Figure 9e). This result suggests that efficiencies can be gained by implementing a hybrid algorithm.

The landscapes associated with three points along the PPF are presented in Figure 10. Points toward the top of the curve are associated with landscapes that are heavily agricultural and, thus, performed well on the gross marginal return objective. Those points toward the bottom of the curve have 25 percent of the landscape in CRP practices that return relatively high EBI values.

Exploring the Solution Space

One benefit of using landscape pattern to calculate the niche count (see Equation 9) is that it promotes geographic diversity within localized regions of the objective space. Two solutions that are close in objective space can represent notably different landscapes, and this spatial diversity provides decision makers with options. An example of such a situation is provided in Figure 11. Though the two solutions shown here lie near one another in objective space, they exhibit different landscape patterns. Solution 10a has fewer hectares in CRP than Solution 10b, but has approximately the same total EBI and income values. The EBI/hectare value for Solution 10a is higher than that of 10b because a larger percentage of its CRP hectares are close to the protected lands of Buttonland Swamp (N1d in the EBI calculation) and a greater percentage of the CRP is enrolled into wetland restoration projects. The land associated with the restored wetlands, however, possesses the most productive soils in the watershed. Solution 10b has more land in crop production than 10a, but since the soil productivity of the tilled land is, on average, lower, it does not generate additional income. Individuals may

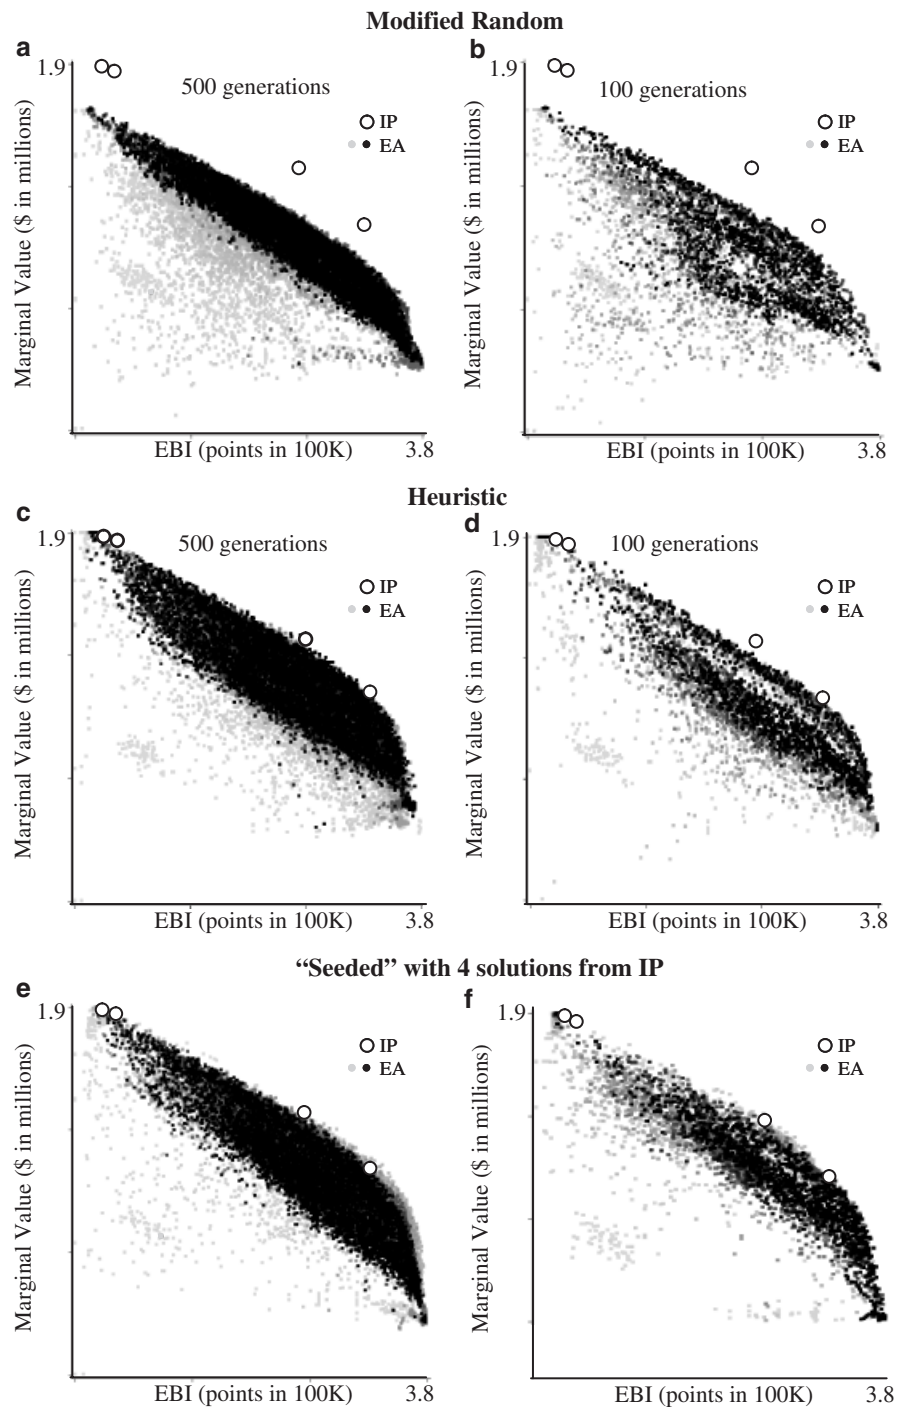


Figure 9. Objective space diagrams produced by the EA. The production possibility frontier was reached using the heuristic initialization algorithm with a population size of 500 and the EA/IP hybrid initialization algorithm with a population size of 100 and 500.

disagree on which solution is preferable, and a final decision may result from compromise.

Finding Compromise Solutions

Because the resolution of semistructured spatial problems often requires compromise and consensus building, tools are needed to expedite the search for alternatives that are satisfactory to all decision makers.

Consider, for example, the following scenario. Decision maker A prefers the alternative represented by Point 1 in Figure 12, perhaps because the spatial pattern of land use is similar to the current situation and because the required level of public investment is relatively low. To reach consensus with other decision makers, however, this individual knows that a compromise alternative that promotes a higher level of environmental quality must be found. How does one go about finding such an

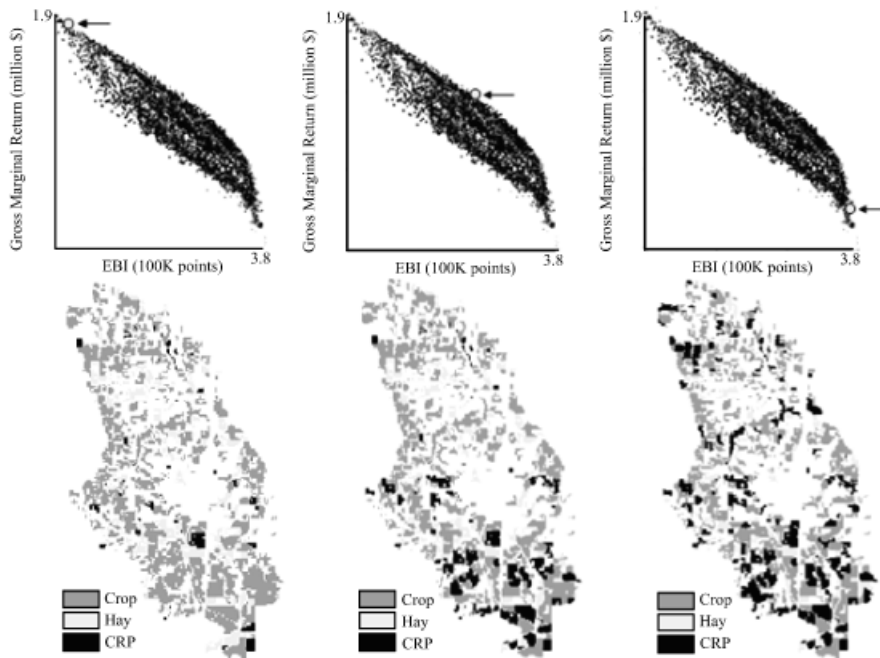


Figure 10. Three representative landscapes illustrating how trade-offs between gross marginal return and EBI become manifest in geographical space.

alternative from among the more than 45,000 alternatives presented in the objective space? For the sake of discussion, assume that a compromise solution must increase EBI scores by at least 20 percent, but not reduce gross marginal return by more than 10 percent. Software was written to support such queries and symbolize the qualifying landscapes as a function of Equation 12 (Figure 12). Using this graphic as a guide, the attention of the decision maker might be led to the solution represented by Point 2, the highest-income producing landscape that meets this criterion. Alternatively, this individual may be curious about the solution represented by Point 3. This point is at the limit of the decision makers' willingness to compromise, but still has a spatial pattern similar to that of Point 1. This situation represents one type of search procedure that may be of interest to decision makers: a search for solutions similar in geographic space to a selected solution, but dissimilar in objective space. The decision makers can just as easily search for solutions similar in objective space, but dissimilar in geographic space.

Reaching for the Frontier

One of the greatest challenges associated with multiobjective evolutionary algorithms is the production of a full and diversified PPF. While many methodologies have been developed to help meet this challenge (see Deb 2000 for a more complete review), none has been implemented in a geographic context. In the earlier section, Fitness Function, we discussed the implementation

of two procedures that are designed to help promote a fully formed curve: niche counting and the use of subpopulations. In this section we investigate the relative impact that these approaches, and associated parameters, have on the development of the front.

The results of this investigation are presented in Figure 13, which shows the effects of using different strategies to produce the PPF. From these illustrations, we see the importance of using both diversification strategies to explore the solution space fully. Figure 13a presents the original solution produced by the heuristic initialization algorithm. Subpopulations alone (Figure 13b) produce a continuous "front" but fail to reach the PPF, devolving back to inferior solutions in the lower portion of the objective space. Niche count alone tends to emphasize compromise solutions in the middle of the PPF (Figure 13c). The use of subpopulations without migration produces specialized phenotypes, and, thus, individuals are clustered into six regions (one for each subpopulation) of the objective space (Figure 13d). High share functions and low population levels can produce reasonable results, but neither quite reaches the frontier (Figures 13e and f). In summary, to reach the PPF it is necessary to:

1. begin with a properly initialized population of sufficient size,
2. instantiate subpopulations and implement inter-population migration, and
3. provide incentives for diversification along the frontier (e.g., implement niche count).

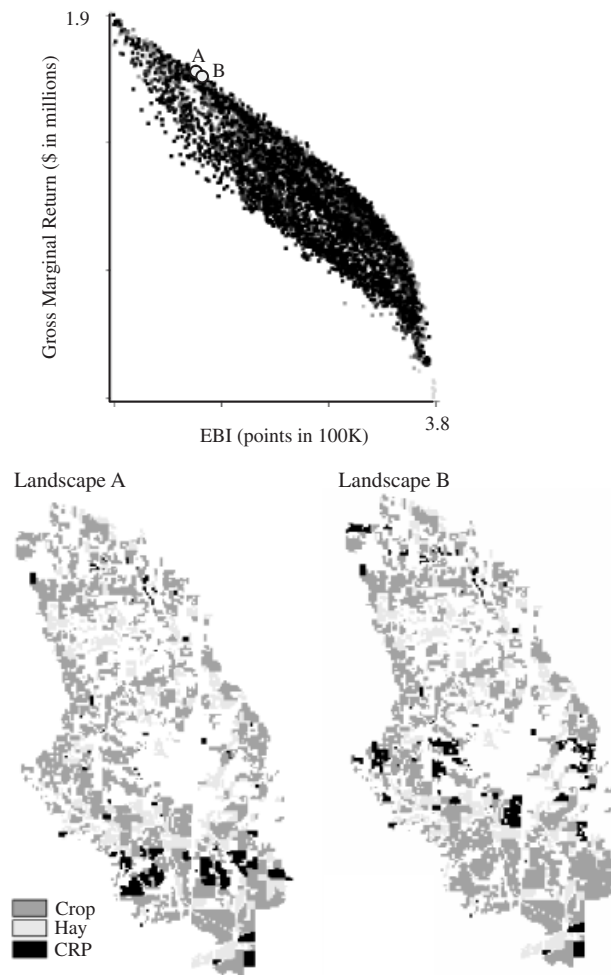


Figure 11. Alternatives close in objective space can have distinctly different spatial patterns.

Discussion

We began this article with the supposition that traditional multiobjective evaluation techniques often cannot present public policymakers with a sufficiently large set of alternative solutions to facilitate well-informed decisions or capture trade-offs among competing objectives. Furthermore, solutions to many real-world geographic problems often prove to be intractable with these techniques. We implemented an EA-based approach that presented to the user 45,000 alternatives distributed within a feasible solution space and across a PPF. Trade-offs, implicit in the form of this graph, can be quantified by fitting a curve through all nondominated points. While few would argue that this approach fails to produce a sufficiently large set of alternative solutions, the size of this set places a significant cognitive burden on policymakers (and analysts) as they try to evaluate the

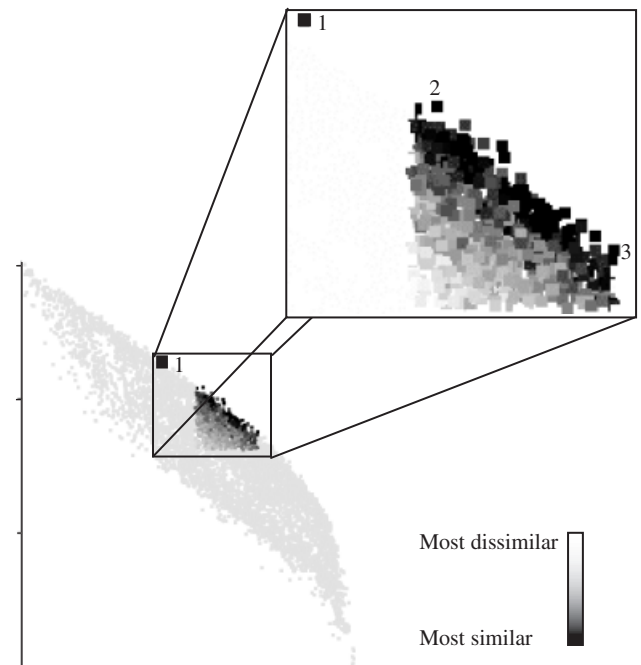


Figure 12. Visualization tools have been developed to help locate interesting or compromise solutions. A decision-maker who prefers Alternative 1, for example, may need to find a compromise solution that has an EBI value at least 20% higher, but does not reduce gross marginal return by more than 10%. The inset illustrates solutions that meet these criteria; darkly shaded points (e.g., Alternative 2) represent solutions that are most like the decision-maker's preferred solution (Alternative 1).

geographic consequences of alternative policies. We attempt to reduce this burden by providing tools to help individuals identify desirable solutions.

Computational intractability often results from the application of spatial constraints to geographic problems. Selecting a subset of sites in a location-allocation problem, searching for an optimal juxtaposition of cover types to support wildlife, and restricting the maximum amount of CRP land in a county, all represent common forms of spatial constraints that lead to intractable problems. As illustrated here, an EA can provide a robust heuristic approach to solving such problems. Though our results indicate that the EA approach shows considerable promise as an analytical method for public policies with geographic consequences, several additional research areas await further examination.

Alternative Objectives

The objectives analyzed as part of this research include agricultural income, environmental quality, and public investment. The EA operates directly on the cover type of farm fields to optimize these three

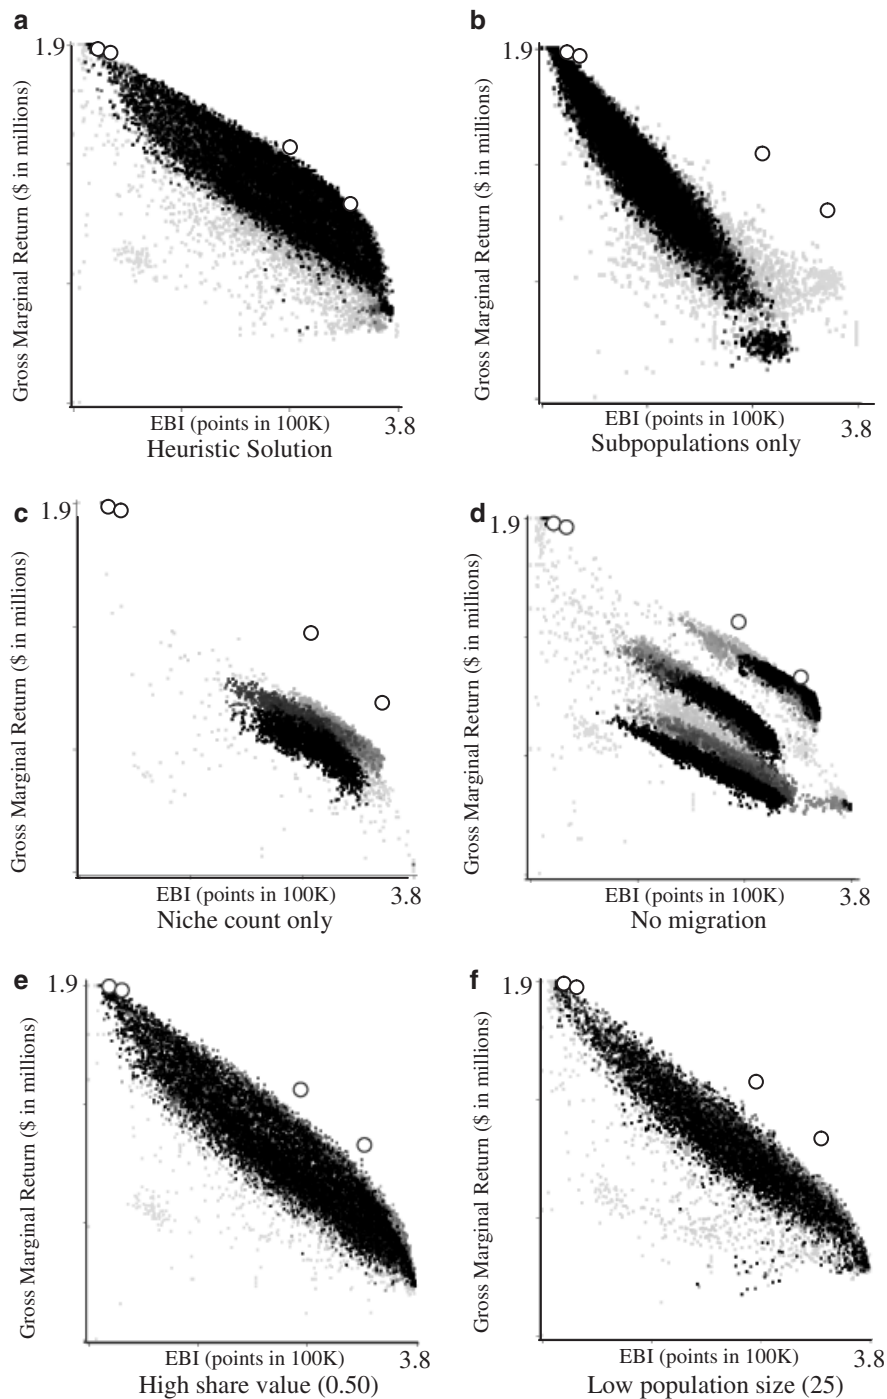


Figure 13. To reach the production possibility frontier using evolutionary algorithms it is important that: (1) the population is properly initialized; (2) the population has a sufficient number of individuals; (3) subpopulations are instantiated; (4) interpopulation migration is allowed; and (5) diversification along the frontier is facilitated (e.g., through implementation of a niche count algorithm).

objectives at the watershed level. While this was viewed as appropriate in the context of public policy, real land-use decisions are usually made at the farm level, and optimizing at the farm level may generate different spatial patterns. Ongoing work that employs multi-agent technology is intended to explore this issue more fully. Furthermore, the EBI was used as a surrogate for environmental quality to mimic existing real-world practice. However, the EBI is in essence a scalarized

multiobjective index that is weighted by point value (e.g., N6 carries 1/4th the weight of N2). Though the relative weight of individual objectives (not all of which are quantitative) is predetermined by law, this index could be “unpacked” to evaluate the effect that the embedded objective weights have on land-use patterns. Alternatively, other indices of environmental quality could be developed and incorporated into this analysis with relative ease.

Enhancements to the Software System

The developed system is an initial effort designed to support future enhancements that would facilitate analyses in other problem domains. Logical extensions include, but are not limited to, modelbase management tools (Bennett 1997) that would allow the user to easily change objectives or EA parameters (e.g., population size, recombination rate, number of generations) and database management tools that would make it easier for users to select only those portions of the decision space that meet their needs.

Validation

While our approach offers considerable promise as a tool that supports the analysis of multiobjective public policies with geographic consequences, we would be remiss if we did not offer some words of caution about validation. Without an independent means of verifying that at least some points on the PPF are found, uncertainty about whether or not the frontier was reached will remain. The graph in Figure 9a, for example, has the appearance of a well-formed PPF, and it is only when we compare these results to those of the IP solutions that we realized that the EA converged prematurely. Once validation points are produced, they can be used to seed the EA and thus promote fast and reliable convergence to the frontier.

Conclusions

A landscape can be managed as a resource that is capable of producing a variety of valuable services. Some of these services can be quantified easily (e.g., crop production), while others are more qualitative (e.g., maintaining a scenic view or type of lifestyle). Decision makers must attempt to produce policies that promote a delicate balance of competing services that meet competing stakeholder needs. Furthermore, decision makers often must meet this complex challenge with limited knowledge about how alternative policies will affect the landscape and the associated trade-offs among competing services. Gaining an understanding about the trade-offs associated with potentially competing services has proven to be difficult because the goal of multiobjective geographic problem solving has often been to identify a small set of plausible solutions. Indeed, the product of such analyses is often a single “best” solution. These approaches may be useful when decision makers can

agree on the relative importance of competing services or when the number of spatial configurations (landscape patterns) that can provide these services is small. When these conditions are not met, however, a more exploratory approach is needed.

In this article we have developed a technique that is designed to produce production possibility frontiers using an EA. These frontiers illustrate the trade-offs that exist among competing services. Each individual in the evolving population represents a landscape tessellated, in this case, into farm fields. The EA modifies land use at the farm field level to optimize three objectives at the watershed level (maximize gross marginal return, maximize environmental benefit, minimize public investment). While the goal is to develop a set of nondominated solutions that define the PPF, the process also produces a large set of near-optimal solutions that might prove useful to decision makers given the semi-structured nature of many geographic problems. The exploration of this solution set and the ultimate selection of a specific solution is facilitated through a user interface that presents the problem from three interrelated perspectives (solution space, objective space, and geographic space).

As we developed this approach, we experimented with three different methods for creating the initial EA population. In the first method, following in the tradition of EA, we created landscapes with random land cover patterns. The second method used farm-field-level indices as heuristics for system level (e.g., watershed level) objectives. Finally, a hybrid EA/IP approach was implemented that used a small number of IP results to “seed” the EA. The results produced by these methods were compared to those obtained using a conventional IP approach. For the presented problem, the populations that were initialized using the heuristic and EA/IP hybrid approaches developed fully formed PPFs. Furthermore, the EA was able to find solutions along portions of the PPF that the IP approach did not. The hybrid EA/IP technique, in particular, seems to have significant potential for the analysis of public policies with geographic consequences. The EA provides the diversity of optimal and near-optimal solutions needed to support semistructured problem solving and the ability to extend analyses into otherwise intractable regions of the decision space, while the IP analysis provides validation and enhanced performance. More generally, this approach makes a significant contribution to geographic information science and public policy analysis because it elucidates a wide range of geographic consequences associated with alternative formulations of public policy.

Acknowledgments

The authors wish to thank the USDA, NSF, and EPA for their joint support of this project through the Water and Watersheds Program, Dr. Steven Kraft and Dr. Jeffery Beaulieu for their advice on issues of agricultural economics, Mr. Girmay Misgna for work on early versions of the EA code, and the reviewers for their thoughtful comments.

References

- Armstrong, M. P., N. Xiao, and D. A. Bennett. 2003. Using genetic algorithms to create multicriteria class intervals for choropleth maps. *Annals of the Association of American Geographers* 93 (3): 595–623.
- Bäck, T. 1996. *Evolutionary algorithms in theory and practice*. Oxford, U.K.: Oxford University Press.
- Bäck, T., U. Hammel, and H.-P. Schwefel. 1997. Evolutionary computation: Comments on the history and current state. *IEEE Transactions on Evolutionary Computation* 1 (1): 3–17.
- Balling, R. J., J. T. Taber, M. Brown, and K. Day. 1999. Multi-objective urban planning using a genetic algorithm. *ASCE Journal of Urban Planning and Development* 12 (2): 86–99.
- Bennett, D. A. 1997. A framework for the integration of geographical information systems and modelbase management. *International Journal of Geographical Information Science* 11 (4): 337–57.
- Bennett, D. A., M. P. Armstrong, and G. A. Wade. 1996. Agent mediated consensus-building for environmental problems: A genetic algorithm approach. *Third International Conference on Environmental Modeling and Geographic Information Systems*, Santa Fe, NM.
- Bennett, D. A., B. Middleton, S. Kraft, C. Lant, R. Sengupta, J. Beaulieu, D. Sharpe, K. Cook, B. Burr, R. Beck, and K. Flanagan. 2001. *Ecosystem function and restoration in the Cache River Bioreserve: Project completion report*. (Carbondale: Southern Illinois University at Carbondale).
- Bennett, D. A., G. A. Wade, and M. P. Armstrong. 1999. Exploring the solution space of semi-structured geographical problems using genetic algorithms. *Transactions in GIS* 3 (1): 51–71.
- Brill, E. D., Jr., J. M. Flach, L. D. Hopkins, and S. Ranjithan. 1990. MGA: A decision support system for complex, incompletely defined problems. *IEEE Transactions on Systems, Man, and Cybernetics* 20 (4): 745–57.
- Brookes, C. J. 2001. A genetic algorithm for designing optimal patch configurations in GIS. *International Journal of Geographical Information Science* 15 (6): 539–59.
- Carver, S. J. 1991. Integrating multi-criteria evaluation with geographical information systems. *International Journal of Geographical Information Systems* 5 (3): 321–39.
- Chambers, L., and M. Taylor. 1996. Planning with GENIE: a model for forecasting and decisionmaking. *Environment and Planning B* 23:697–709.
- Coello, C. A. C. 2000. An updated survey of GA-based multiobjective optimization techniques. *ACM Computing Surveys* 32 (2): 109–43.
- Cohon, J. L. 1978. *Multiobjective programming and planning*. New York: Academic Press.
- Deb, K. 2001. *Multi-objective optimization using evolutionary algorithms*. Chichester, U.K.: John Wiley & Sons.
- Densham, P. J. 1991. Spatial decision support systems. In *Geographical information systems: Principles and applications*, ed. D. J. Maguire, M. F. Goodchild, and D. W. Rhind, 403–12. Harlow, U.K.: Longman Scientific and Technical.
- Dibble, C., and P. J. Densham. 1993. Generating interesting alternatives in GIS and SDSS using genetic algorithms. In *Proceedings of GIS/LIS '93*, 180–9. Minneapolis, MN: ACSM, ASPRS, AM/FM International, AAG, URISA.
- Farmdoc 2003. <http://www.farmdoc.uiuc.edu/> (last accessed 30 July 2004).
- Fogel, L. J. 1962. Autonomous automata. *Industrial Research* 4:14–19.
- Fonseca, C. M., and P. J. Fleming. 1993. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In *Proceedings of the Fifth International Conference on Genetic Algorithms*, ed. S. Forrest, 416–23. San Mateo, CA: Morgan Kaufmann.
- . 1995. An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation* 3 (1): 1–16.
- Goldberg, D. E. 1989. *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley.
- Holland, J. H. 1975. *Adaptations in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Hosage, C. M., and M. F. Goodchild. 1986. Discrete space location-allocation solutions from genetic algorithms. *Annals of Operations Research* 6:35–46.
- Jankowski, P. 1995. Integrating geographical information systems and multiple criteria decision making. *International Journal of Geographical Information Systems* 9 (3): 251–73.
- Koza, J. R. 1992. *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: The MIT Press.
- Krzanowski, R. M., and J. Raper. 1999. Hybrid genetic algorithm for transmitter location in wireless networks. *Computers, Environment and Urban Systems* 23:359–82.
- Malczewski, J. 1999. *GIS and multicriteria decision analysis*. New York: John Wiley & Sons.
- Matthews, K. B., A. R. Sibbald, and S. Craw. 1999. 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.
- Miettinen, K. M. 1999. *Nonlinear multiobjective optimization*. Boston: Kluwer Academic Publishers.
- Mimouni, M., S. Zekri, and G. Flichman. 2000. Modelling the trade-offs between farm income and the reduction of erosion and nitrate pollution. *Annals of Operations Research* 94:91–103.
- Rechenberg, I. 1965. *Cybernetic solution path of an experimental, problem library translation No. 1122*: Royal Aircraft Establishment.
- Ribaudo, M. O., D. L. Hoag, M. E. Smith, and R. Heimlich. 2001. Environmental indices and the politics of the conservation reserve program. *Ecological Indicators* 1 (1): 11–20.
- Rittel, H. W. S., and M. M. Weber. 1973. Dilemmas in a general theory of planning. *Policy Sciences* 4:155–69.
- Rupe, J. C., R. T. Robbins, and E. E. Gbur. 1997. Effects of crop rotation on soil population densities of *Fusarium solani* and *Heterodera glycines* and on the development of sudden death syndrome of soybeans. *Crop Protection* 16 (6): 575–80.

- Saaty, T. L. 1980. *The analytical hierarchy process*. New York: McGraw Hill.
- Samet, H. 1990. *The design and analysis of spatial data structures*. Reading MA: Addison-Wesley.
- Schwartz, B. 2004. *The paradox of choice: Why less is more*. New York: Ecco Press.
- Sprague, R. H., Jr., and E. D. Carlson. 1982. *Building effective decision support systems*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Srinivas, N., and K. Deb. 1994. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation* 2 (3): 221–48.
- Xiao, N., and M. P. Armstrong. 2003. A specialized island model and its application in multiobjective optimization. In *Genetic and evolutionary computation GECCO 2003*, ed. E. Santú-Paz, J. A. Foster, K. Deb, L. D. Davis, R. Roy, U.-M. O'Reilly, H.-G. Beyer, R. Standish, G. Kendall, S. Wilson, M. Harman, J. Wegener, D. Dasgupta, M. A. Potter, A. C. Schultz, K. A. Dowsland, N. Jonoska, and J. Miller, 1530–40. Heidelberg: Springer-Verlag.
- Xiao, N., M. P. Armstrong, and D. A. Bennett. 2002. ChoroWare: a software toolkit for choropleth map classification. In *New tools for spatial data analysis: Proceedings of CSISS Specialist Meeting*, ed. L. Anselin and S. J. Rey. CD-ROM. Santa Barbara, CA: The Center for Spatially Integrated Social Science.
- Xiao, N., D. A. Bennett, and M. P. Armstrong. 2000. Solving spatio-temporal optimization problems with genetic algorithms: A case study of a seed dispersal model. In *Proceedings of the 4th International Conference on Integrating GIS and Environmental Modeling: Problems, prospects, and needs for research*, ed. B. O. Parks, K. M. Clarke, and M. P. Crane. Banff, Canada.
- . 2002. Using evolutionary algorithms to generate alternatives for multiobjective site search problems. *Environment and Planning A* 34 (4): 639–56.

Correspondence: Department of Geography, The University of Iowa, 316 Jessup Hall, Iowa City, IA 52242, e-mail: david-bennett@uiowa.edu, (Bennett); marc-armstrong@uiowa.edu (Armstrong); Department of Geography, 1132 Derby Hall, 154 North Oval Mall, Columbus, OH 43210, e-mail: xiao.37@osu.edu (Xiao).