A Multiple-Paradigm System for Rangeland Pest Management

John D. Hastings and L. Karl Branting
Department of Computer Science
University of Wyoming
Laramie, Wyoming 82071-3682

Jeffrey A. Lockwood
Department of Plant, Soil, and Insect Sciences
University of Wyoming
Laramie, Wyoming 82071-3354

Abstract

Polycultural agroecosystems, such as rangelands, are too complex and poorly understood to permit precise numerical simulation. Management decisions that depend on behavioral predictions of such ecosystems therefore require a variety of knowledge sources and reasoning techniques. Our approach to designing a computer system that provides advice concerning such ecosystems is to incorporate various reasoning paradigms and apply whatever paradigm is most appropriate to each task arising in the advice process. This approach is based on a particular process description of expert human problem solving that uses four different reasoning paradigms: model-based reasoning; case-based reasoning; rule-based reasoning; and statistical reasoning. The process description is implemented in CARMA, a computer system for advising ranchers about the best response to rangeland grasshopper infestations. CARMA attempts to emulate the human ability to integrate multiple knowledge sources and reasoning techniques in a flexible and opportunistic fashion. The goal of this approach is to enable computer systems to optimize the use of the diverse and incomplete knowledge sources and to produce patterns of reasoning that resemble those of human decision-makers.

Introduction

One of the most striking characteristics of human problem-solving behavior is the ability to exploit multiple knowledge sources and reasoning techniques. This ability is important because problem solving often occurs in an environment of incomplete knowledge. Automating this ability requires techniques for integrating multiple problem-solving paradigms in a flexible manner.

For example, many types of monitoring, diagnostic, and planning tasks require predicting the behavior of a physical system. Highly accurate models exist for the behavior of many simple physical systems. However, models of complex biological and ecological systems are often incomplete, either because a complete state description for such systems cannot be determined or because the number and type of interactions between system elements are poorly understood. Under these circumstances, other knowledge sources and reasoning techniques, such as case-based reasoning and statistical reasoning, are essential for prediction.

Agriculture is moving towards complex polycultures of organisms to enhance the economic and environmental sustainability of food and fiber production (Edwards *et al.* 1990; Miller 1992; Chadwick & Marsh 1993). As we diversify our agroecosystems, we will be increasingly confronted with novel, complex management problems for which we have very little data and only limited experience. Fully parameterized, accurate simulation models that were possible with classically simplified agricultural operations (*i.e.*, monocultures of annual plants) are not a viable approach to highly complex systems (Pimm 1991; Allen & Hoekstra 1992; Lockwood 1994).

The complexities of polycultural agroecosystems are typified by rangelands. Rangelands are diverse communities of mixed annual and perennial plants that are highly integrated with livestock and wildlife production. Rangeland ecosystems present an abundance of critical management problems, including optimal stocking rates and grazing systems, water development, wildlife enhancement, noxious weed control, and insect pest management. The study and practice of rangeland management over many decades has led to the development of perhaps the richest knowledge base for any complex, integrated, temperate agroecosystem. However, the complexity of these systems places the quantitative data necessary to develop and parameterize simulation models for use by managers well beyond our current, or foreseeable, capacity (Lockwood & Lockwood 1991; 1993). This is not to say that simulation models are not effective research tools. Rather, these models do not take advantage of the considerable qualitative experience that characterizes the knowledge base. Moreover, there are important management decisions that must be made in the context of significant quantitative uncertainty.

This paper describes an artificial intelligence approach to agricultural management that integrates a variety of reasoning paradigms. Underlying this approach is the view that computer systems should emulate the human capacity to employ whatever reasoning technique is most appropriate for a given task or knowledge type. The goal of this approach is to enable computer systems to optimize the use of the diverse and incomplete knowledge sources available to agricultural decision-makers and to produce patterns of reasoning that more closely resemble those of human decision-makers.

The next section of this paper describes the task of rangeland grasshopper pest management, argues that model-based reasoning standing alone is insufficient for this task, and sets forth a process description of expert human problem solving for this task that uses four different reasoning paradigms: model-based reasoning; case-based reasoning; rule-based reasoning; and statistical reasoning. CARMA, a computer system that models human use of multiple reasoning paradigms for rangeland grasshopper pest management, is described in the following section. The final section discusses the relationship of CARMA to other research on: (1) the application of artificial intelligence to natural resources problem-solving, and (2) the integration of multiple artificial intelligence techniques.

An Information Processing View of Rangeland Grasshopper Pest Management

In most years and locations, the majority of grasshopper species are innocuous or even beneficial to grassland ecosystems. Of over 300 species of grasshoppers in the western United States, perhaps only 15 can be considered serious pests; many of the other species are beneficial in terms of controlling weeds, nutrient cycling, and food for wildlife (Lockwood 1993a; Lockwood 1993b). However, large-scale grasshopper outbreaks are capable of inflicting serious economic damage to western livestock producers. On average, grasshoppers annually consume 21-23% of western rangeland forage, at an estimated loss of $400 million (Hewitt & Onsager 1983). For example, in 1985-86, Wyoming treated approximately 6.5 million acres at a cost of $22.75 million to private, state and federal interests. Due in part to a restructuring of the state cost-share program, some 9 million acres of infested rangeland were left untreated in 1987, resulting in the loss of 225,000 tons of air-dried forage. The decision whether to use insecticides or other control measures is a complex task because of the multiplicity of relevant factors, such as maintaining minimal inputs for profitable ranching in the western United States, preserving natural enemies for chronic control of grasshoppers (Joern & Gaines 1990), safeguarding biodiversity, and protecting environmental and human health.

Management of rangeland grasshoppers requires a variety of knowledge sources. Because of the diversity of rangeland conditions, productivities, and vegetation types, the enormous range of weather conditions, and site-specific elements (*e.g.*, honey production, wildlife management, water developments, etc.) the decisions of pest managers addressing this task appear to us to be largely based on a set of synthetic, prototypical cases. These cases do not necessarily correspond to specific real world experiences, but express the essential features of past management experiences that define prototypical instances in which particular management practices are optimized. For example, the most authoritative guide for management of African (Desert) Locust infestations consists of a collection of discrete cases compiled as a reference for workers (Pedgley 1981)[1]. Specific cases have been found to be an important knowledge source in a variety of problem domains in which precise general rules are unavailable or inadequate (Klein & Calderwood 1988).

A second form of knowledge used by experts in rangeland grasshopper management consists of general rules that appear to constrain management decisions under specific circumstances (*e.g.*, there is no point in controlling grasshoppers once the adults have laid eggs, as the majority of the damage is already done and the next generation is assured). Finally, there are some aspects of grasshopper population ecology for which we have sufficient information to apply mechanistic models (*e.g.*, the development of grasshoppers through the nymphal stage is a relatively well-defined function of temperature).

While model-based reasoning can play a role in grasshopper management, there is a general recognition that the interactions affecting grasshopper population dynamics are too poorly understood and too complex to permit precise prediction through numerical simulation (Lockwood & Lockwood 1991; Pimm 1991; Allen & Hoekstra 1992). Grasshopper populations are extremely labile, and a multiplicity of biotic and abiotic factors regulate their densities. Based on a simplified rangeland habitat comprised of just 10 grasshopper species, 10 plant species, four soil types, and 10 predators, Lockwood (1994) calculated that 10,560 two- and three-factor interactions end with grasshoppers. More realistic estimates of diversity suggest as many as 175 million interactions end with grasshoppers. If even 1% of these were ecologically relevant, the number of interactions would be far too great to simulate. Thus, while simulations of grassland ecosystems can provide insight into their dynamics (Fedra 1991; Rodell 1978), it is not feasible to devise models adequate for accurate prediction of the consequences of treatment options.

However, the ability of entomologists and pest managers to provide meaningful advice indicates that other sources of knowledge can compensate for the absence of a complete model of rangeland ecosystems. To explicate these knowledge sources, we performed a protocol analysis of problem solving by several experts in rangeland grasshopper management at the University of Wyoming. For each expert, we transcribed several problem-solving

---

[1]Although not explicitly setting forth any systematic or formal means of case-matching, Pegley (1981) is premised on the assumption that forecasting locust population dynamics and appropriate management strategies can be based on comparisons with specific cases.

episodes in which the expert responded to a simulated telephone inquiry by a rancher. These "solve-aloud" problem-solving episodes illustrate the elicitation of relevant case facts by the expert, the formation and discrimination among tentative hypotheses, and expert explanations. Based on this protocol analysis, we have developed the following process description of expert problem solving for this task:

1. Determine the relevant facts of the infestation case, such as grasshopper species, developmental periods[2], and density, from information provided by the user. This requires **rule-based reasoning** using rules such as, "if brightly colored wings are observed in the spring, then the grasshoppers are adults that overwintered as nymphs."

2. Determine whether the grasshopper infestation is a potential problem. The infestation is not a problem if:

    a. The current date is outside of the "critical period" when forage needed for livestock grows. This requires **rule-based reasoning** to determine whether the date is outside of the critical period, given the historical critical forage growing period for the location and the date.

    b. The size of the infestation is small. This requires **rule-based reasoning** to determine whether the infestation size is below a minimum threshold.

    c. The majority of the grasshoppers overwinter as nymphs. This requires **model-based reasoning** to estimate whether the nymphal and adult periods of the current grasshoppers occur when only nymphal-overwintering grasshoppers are possible, given the grasshoppers' developmental periods, the growing season dates for the location, and the date.

    d. The majority of the grasshoppers are in inappropriate periods. This requires **rule-based reasoning** to determine whether the majority of the grasshoppers are at such an early developmental period that the extent of the infestation cannot be predicted with reasonable certainty or at such a late developmental period that a significant proportion of lifetime forage consumption

[2]During their lifetime, grasshoppers progress through three developmental stages: egg, nymph, and adult. The nymphal stage usually consists of five instars separated by molts. We define the **developmental periods** of a grasshopper's lifecycle to include egg, five nymphal instars, and adult.

and egg-laying have already occurred, making insecticide application pointless.

3. If the infestation is potentially a problem, determine whether grasshopper consumption will lead to competition with livestock for available forage.

   a. Estimate the proportion of available forage that will be consumed by each distinct grasshopper population (*i.e.*, nymphal overwintering, egg overwintering) using **case-based reasoning** For each distinct grasshopper population (*i.e.*, subcase):

      i. Determine the prototypical case that most closely matches the current subcase. This requires **model-based reasoning** to assist matching by aligning the developmental periods of the prototypical case and the subcase.

      ii. Adapt the consumption estimate predicted by the prototypical case based on the featural differences between the prototypical case and current subcase. This requires **model-based reasoning** to account for the influence of each feature on consumption.

      iii. Further adapt the consumption estimate so that it accounts only for "critical period" consumption. This requires **model-based reasoning** to determine the proportion of lifetime consumption occurring during the critical period for both the prototypical case and subcase given their developmental periods and the date.

   b. Total the forage loss estimates for each subcase to predict the overall proportion of available forage that will be consumed by grasshoppers.

   c. Compare grasshopper consumption with the proportion of available forage needed by livestock.

4. If there will be competition, determine what possible treatment options should be excluded. This requires **rule-based reasoning** using rules such as "wet conditions preclude the use of malathion; environmental sensitivity precludes all chemical treatments."

5. If there are possible treatment options, for each one provide an economic analysis by estimating both the first-year and long-term savings.

a. Estimate the first-year savings using **model-based reasoning** to determine the proportion of forage which would be saved given the efficacy of the treatment type, the developmental periods of the grasshoppers at the time of treatment, and the proportion of lifetime consumption by grasshoppers at each period.

b. Estimate the long-term savings using **rule-based reasoning** to determine if the majority of the grasshoppers will begin laying eggs before treatment can be applied given the developmental distribution of the grasshoppers at the time of treatment. If the majority of grasshoppers will not begin laying eggs, use **statistical reasoning** to determine the decreased probability of infestation in subsequent years given the Markov transitional probabilities for the infestation location and the efficacy of the treatment type.

In performing this process, human experts exhibit several characteristics that would be desirable to emulate in a computer system:

- **Speed.** Human experts can provide useful advice very quickly. This suggests, consistent with our process description, that human experts can use highly compiled knowledge in the form of prototypical cases and rules.

- **Graceful degradation.** Human experts can use, but do not require, highly precise information of the type required for accurate model-based reasoning. Less accurate information may degrade the quality of advice an expert can give, but doesn't preclude useful advice. In the worst case, human experts can provide plausible advice based merely on the location of the rangeland and the date.

- **Explanations in terms of a causal model.** Although the speed and graceful degradation of human expert performance suggest that experts can use compiled knowledge, they can also readily provide causal explanations for their conclusions. Moreover, entomologists can generate causal predictions of the effects of incremental variations on case facts. This behavior strongly suggests that they have access to causal models that can assist in explanation and in adaptation of prototypical cases.

- **Opportunism.** Human experts can use a variety of different strategies to solve a single given problem depending on the available information. Human experts don't address the subgoals that arise in decision-making in an invariant order, but adapt their problem-solving behavior to the particular facts of a given case.

In summary, rangeland grasshopper management typifies a task in which the absence of a complete and accurate model necessitates integrating a variety of individually incomplete knowledge sources. The next section describes CARMA, a multiple-paradigm computer system for rangeland management.

CARMA: A Multiple-Paradigm Reasoner

CARMA (CAse-based Range Management Adviser) is a system for advising ranchers about the best response to rangeland grasshopper infestations. CARMA implements the process description of entomological problem solving set forth in the previous section.

CARMA was originally implemented as a purely case-based reasoning system using the PROTOS case-based reasoning shell (Porter, Bareiss, & Holte 1990). However, this implementation proved to be a poor model of expert problem solving in this domain. PROTOS is designed to produce a diagnostic category as a solution. However, our protocol analysis indicated that a solution should consist of a treatment recommendation supported by an explanation in terms of causal, economic, and pragmatic factors, including a numerical estimate of the proportion of forage consumed and a cost-benefit analysis of the various treatment options. The rule-based and model-based steps of expert problem solving set forth in the previous section, which are necessary for such solutions of this nature, can't be accommodated within a purely case-based reasoning approach. The focus of the CARMA project has therefore been on integrating the multiple problem solving paradigms used by human experts.

Figure 1 shows an overview of CARMA's components, including the consultation manager and its tasks, and the reasoning modules and information required to complete the tasks. The following sections describe CARMA and its use of different reasoning paradigms.

Determining Relevant Case Features

CARMA provides advice by reasoning about the relevant features of an infestation case (*e.g.*, the species, density, and developmental periods of the grasshoppers). These features are inferred by rules from information provided by the user through window-based interface procedures. CARMA makes use of multiple levels of rules for inferring each case feature, ordered by the certainty or the accuracy of each rule. The rules are applied in succession until either the user can provide the necessary information or a default value is chosen. For example, if the value of the case feature "total number of grasshoppers per square yard" is unknown to the user, CARMA attempts to infer this feature using a rule that grasshopper density is equal to 1.5 times the number of grasshoppers seen with each step taken by the user in the field. If the user can't provide the information required for this rule, the system uses the historic average for the area. By applying rules in the order of their certainty, CARMA reasons with the best information available.

A typical interface window for the area of infestation appears in Figure 2. It includes the options "Why" for describing why particular information is important to the consultation, "Help" for advising the user about the various window features and their operations, "Not sure" to trigger the selection of an alternative rule for inferring the feature, and "Done" to indicate that the user has chosen an answer.

Figure 3 shows an input window that asks the user to provide the infestation location by clicking on a map of Wyoming's major roads, towns, and county borders. CARMA uses this location to retrieve the historical values for the site including infestation history, range value, temperature, and precipitation.

Since a complete case specification is not always required to give advice, CARMA fills in the facts of a new case opportunistically. This means that CARMA asks the user for information only when the corresponding case feature is required for the reasoning process to continue. At the earliest point at which a decision can be made, the case-feature inference process halts, advice is given, and the consultation is completed. This minimizes the amount of input required for CARMA to make a decision, thereby accelerating consultations. For example, the date and location of an infestation may indicate that it is too early to assess the severity of a grasshopper infestation. In such cases, CARMA advises the user to rerun the consultation at a later time without prompting for further information.

Determining Infestation Potential

CARMA's first step in advising a rancher about the best response to a grasshopper infestation is deciding whether a potential problem exists. CARMA determines that an infestation is not a problem and terminates a consultation if it discovers any of the following facts:

1. The date is outside of the "critical period".

2. The size of the infestation is too small to be viable.

3. The majority of the grasshoppers overwinter as nymphs.

4. The developmental period of the majority of the grasshoppers is too early for accurate prediction or too late for effective treatment.

The sections that follow describe why these conditions lead to the termination of a consultation, and the methods used to make these decisions.

Outside of the Critical Period

Consumption by grasshoppers is only damaging if it occurs during the critical period of rangeland vegetation when forage needed for livestock grows. If the date of the current infestation is significantly earlier or later than the historical critical period for the area, any grasshoppers that are present will not cause appreciable damage, so no action should be taken. This decision is made by comparing the current date with the historical critical period for the area.

Small Infestation Size

The size of an infestation is an indicator of its viability and hence its future damage potential. A small infestation size indicates either an isolated hatching area or a population with little viability. Infestations smaller than 500 acres are considered unlikely to lead to a significant infestation.

Grasshoppers Overwintering as Nymphs

A tract of rangeland invariably contains multiple grasshopper species. Although virtually all species have only one generation per year, the timing of life history events and consumption characteristics vary greatly. Specifically, grasshoppers overwintering as nymphs divide their consumption between two growing seasons and consume far less during the critical period than grasshoppers overwintering as eggs. If the majority of the grasshoppers have the former life history, little forage loss will occur unless densities are extraordinarily high.

To determine the proportion of nymphal-overwintering grasshoppers, CARMA factors the overall population of a case into subcases according to life history. CARMA initially splits the overall population into observed categories (bandwinged, spurthroated, and slantfaced) based on observed physical characteristics, then classifies grasshoppers into the appropriate overwintering categories (nymphal or egg) based on the typical life histories of the observed categories (*e.g.*, adult bandwinged grasshoppers seen in the spring overwintered as nymphs). For example, the new case set forth in Table 1 is split into two subcases, SubcaseA and SubcaseB, based on life history.[3] The user is given a warning message if the observed population distribution is unlikely given the time of year. For example, if the observed population distribution contains adult bandwinged grasshoppers after August 1, CARMA will warn the user that under normal circumstances only nymphal bandwinged grasshoppers are expected.

Inappropriate Grasshopper Periods

To provide a meaningful consultation, the majority of grasshoppers must be sufficiently developed to determine the extent of the infestation with some certainty (the infestation potential of a very young population of grasshoppers can fluctuate drastically based on the weather and disease and is therefore much less predictable), but immature enough that a significant proportion of lifetime forage consumption remains, making insecticide application or biological control economically feasible. Rule-based reasoning is used to end consultations involving grasshopper populations whose average developmental period is less than nymphal instar 1.5 (*i.e.*, too early) or equal to "Adult" (*i.e.*, too late).

Determining Forage Competition

If a grasshopper infestation is potentially a problem, CARMA estimates forage consumption using a library of prototypical cases. This forage consumption estimate is used to predict whether grasshopper consumption will lead to competition with livestock for available forage.

Prototypical Infestation Cases

Since our protocol analysis indicated that pest managers estimate forage consumption by comparing new cases to prototypical infestation scenarios, we elicited a set of prototypical cases from an entomologist experienced at the rangeland grasshopper pest management advising task who participated in the protocol

[3]This case is used merely to illustrate several of CARMA's features. In a typical consultation, a case occurring so late in the growing season would be classified as "too late" for purposes of insecticide application. A complete analysis of such a case would proceed only if requested by the user.

analysis. An initial set of prototypical cases was obtained by asking the expert what stereotypical situations were used as a standard for comparison with the problem situations addressed by the expert in the protocol analysis. Additional prototypical cases were obtained by presenting the expert with a wide range of artificial problems and asking the expert to identify stereotypical situations that would be most relevant to forage consumption predictions in those situations.

The prototypical cases differ significantly from conventional cases. First, the prototypical cases are not expressed in terms of observable features (*e.g.*, "Whenever I take a step, I see 4 or 5 grasshoppers with brightly colored wings fly"), but rather in terms of abstract derived features (*e.g.*, "A low or moderate density of nymphal-overwintering grasshoppers in the adult period"). In addition, the prototypical cases are extended in time, representing the history of a particular grasshopper population over its lifespan.

CARMA's case library currently consists of eight prototypical cases. Each prototypical case is represented by a "snapshot" at a particular, representative point in time selected by the entomologist. In general, this representative point is one at which the grasshoppers are at appropriate periods (as described above). An example prototypical case appears as Case8 in Table 1.

Case-Based Prediction of Forage Consumption

As previously mentioned, CARMA uses case factoring to split the overall population of a case into subcases according to life history (*i.e.*, overwintering as nymphs or eggs). To predict overall forage loss, CARMA totals forage loss predictions for each subcase. The following sections detail how CARMA predicts forage loss by using a causal model to assist case-based reasoning in three different ways: temporal projection; featural adaptation; and critical period adjustment.

Temporal Projection To predict the forage loss of a subcase, CARMA first retrieves all prototypical cases whose life history matches that of the subcase. Since prototypical cases are extended in time but are represented at a particular time, matching requires temporally projecting the prototypical cases forwards or backwards to align the average developmental period of the new subcase. This requires using the model to simulate grasshopper attrition, which depends on developmental period, precipitation, and developmental rate (which in turn depends on temperature) throughout the interval of the projection. An example appears in Figure 4.

The projected prototypical case whose weighted featural difference from the new case is least is selected as the best match. For example, the prototypical case that best matches SubcaseA after projection is Case8, as shown in Table 1. Because the developmental period of Case8 before projection is later than that of SubcaseA, Case8 must be projected backwards in time, causing grasshoppers that had been lost to attrition to be added back to the population.

Temporal projection aligns developmental periods but not necessarily dates. For example, the date of Case8 after projection is later than the date of SubcaseA because the hatch date of Case8 was later than that of SubcaseA. As a result, the average developmental period of the grasshoppers in SubcaseA on August 20 is the same as that of Case8 two weeks later on September 2.

Featural Adaptation The consumption predicted by the best matching prototypical case is modified to account for any featural differences between it and the subcase. This adaptation is based on the influence of each feature on consumption as represented by global feature weights. For example, a lower temperature value means lower forage losses, because lower temperatures tend to slow development, increasing grasshopper attrition. Thus, the forage loss estimate predicted by Case8 - 15% - must be adapted downward somewhat to account for the fact that temperatures in SubcaseA (cool) are lower than in Case8 (normal).

Feature weights are set using a hill-climbing algorithm that optimizes CARMA's predictive accuracy on training instances. The weights used in featural adaptation constitute a linear approximation of the function from derived case features to consumption amounts in the neighborhood of each prototypical case.

Critical Period Adjustment As previously mentioned, consumption is only damaging if it occurs during the critical forage growing period of a rangeland habitat. The forage loss predicted by a prototypical case must be modified if the proportion of the lifespan of the grasshoppers overlapping the critical period differs significantly between the new case and the prototypical case. This process, termed *critical period adjustment*, requires determining the developmental periods of the new and prototypical cases that fall within the critical period and the proportion of lifetime consumption occurring in these developmental periods.

An example of critical period adjustment appears in Figure 5. Because grasshopper development in SubcaseA is ahead of that in Case8 (SubcaseA's developmental period on August 20 corresponds to Case8's developmental period on September 2), CARMA determines that SubcaseA applies to more of the critical period than Case8 because it will reach instar 5 by the end of the critical period, while Case8 will only reach instar 2. CARMA uses a model of grasshopper's rate of consumption at each developmental period to calculate the proportion of lifetime consumption occurring before the end of the critical period. For example, 47% of SubcaseA's consumption occurs during the critical period, whereas only 6% of Case8's consumption occurs within this period. CARMA therefore scales the initial consumption estimate by (47 ÷ 6) = 7.8.

After adaptation, the consumption predictions for each subcase (*i.e.*, populations of grasshoppers with distinct feeding patterns) are summed to produce an overall consumption estimate. In the given case, the sum of predicted consumption of the two subcases is 57%. Because of the variability resulting from the imprecise nature of rangeland ecosystems, this prediction is converted to the qualitative range, **high-moderate**, meaning that approximately 40 to 60% of the available forage will be lost.

If the proportion of available forage that will be lost to grasshoppers and the proportion needed for livestock (and wildlife) exceeds 100% of the forage available, CARMA concludes that competition will occur. In this example, competition is possible and the consultation should continue if the proportion of available forage needed by livestock is greater than 40%. For example, if forage need is 60%, the expected year-long competition should range from 0% (*i.e.*, (40+60)-100) to 20% (*i.e.*, (60+60)-100). A typical interface window showing forage loss and asking for forage need is shown in Figure 6.

Determining Treatment Options

If there will be competition, CARMA applies a set of rules to determine what possible treatment options are excluded by the conditions of the case. Some of the information necessary for determining exclusion is already known from the case features (*e.g.*, the presence of grasshoppers in the first nymphal instar suggests an ongoing hatch, thereby excluding malathion and carbaryl bait from consideration). Other conditions must be determined from further user input (*e.g.*, "Will it be wet at the time of treatment?" If so, exclude malathion.).

Economic Analysis

For each possible treatment option, CARMA provides estimates of first-year and long-term savings. A typical economic analysis of possible treatment types appears in Figure 7. Notice that this analysis includes "no treatment" as an option. Each analysis involves a range that indicates best to worst case estimates (negative values indicate a loss).

First-year Savings

For each possible treatment option, CARMA estimates the first-year savings as the difference between the value of forage in competition saved by treating and the treatment cost. An example appears in Table 2.

CARMA first computes the amount of pre-treatment forage loss. This is done by temporally projecting the developmental distribution of each subcase forwards to the user-provided treatment date (often a week or more from the current date). In a manner similar to critical period adjustment, CARMA applies a model of grasshoppers' rate of consumption at each developmental period to each subcase to calculate the proportion of lifetime consumption occurring before the treatment date. This proportion is used to scale the year-long forage loss estimate, resulting in the pre-treatment loss. The pre-treatment forage loss estimates for each subcase are summed to produce the total pre-treatment forage loss. Next, CARMA estimates the amount of post-treatment forage loss without treatment by subtracting pre-treatment forage loss from total forage loss. For example, if total forage loss is estimated to be 40-60%, and pre-treatment forage loss is estimated to be 20-30% (*i.e.*, 50% of the grasshoppers' total consumption will occur before the treatment date), then the post-treatment forage loss will be 20-30% (because 50% of the grasshoppers' lifetime consumption must occur after the treatment date).

For each option, CARMA estimates the amount of post-treatment forage loss with treatment according to the expected efficacy of the treatment and the post-treatment forage loss without treatment. For example, the insecticide malathion is usually 80 to 90% effective. If the estimated post-treatment forage loss without treatment is 20-30%, then at best malathion should prevent 90% of the 20% loss, and at worst prevent 80% of the 30% loss, resulting in a 2 to 6% post-treatment forage loss.

CARMA calculates the year-long forage loss for each option by summing pre- and post-treatment forage loss. Year-long competition resulting from an option is calculated by comparing year-long forage loss resulting from the option to forage need. The proportion of forage in competition that will be saved is simply the proportion of forage in competition without treatment minus the proportion of forage in competition with treatment. For example, if pre-treatment forage loss is 20-30% and post-treatment forage loss is 2-6%, the year-long forage loss for the option is 22-36%. Given a forage need of 60%, the year-long competition with treatment ranges from (22+60)-100 = -12 to (36+60)-100 = -4, which is less than zero, resulting in no competition. If the year-long forage in competition without treatment is 0-20%, and the treatment option will result in no competition, then the expected forage in competition saved by treating is 0-20%.

With a per-unit forage value estimate provided by the user, CARMA estimates the first-year savings for an option to be the value of forage in competition that is saved minus the cost of the treatment.

Long-term Savings

CARMA uses statistical reasoning and the historically derived Markov transitional probabilities for the infestation location to calculate long-term savings. First, CARMA determines whether the grasshoppers will begin laying eggs before the treatment date. If the developmental distribution of the grasshoppers at treatment is dominated by adults, CARMA determines that too many eggs will already be laid, and no long-term savings will result from treatment because eggs are not affected by treatment.

If few eggs will have been laid, CARMA calculates the savings for future years for each treatment type as the value of year-long (*i.e.*, total) forage in competition without treatment (taken from the first year calculations) times the difference between the probability of infestation without treatment and the probability of infestation with treatment. The transitional probabilities are adjusted based on the expected efficacy of the treatment. Future savings are calculated for each year as long as the probability of infestation with treatment for the year is significantly lower that the probability of infestation without treatment (*i.e.*, until the benefits of treatment have ended). Following the economic analysis, the consultation is complete.

Multiple-Paradigm Reasoning in CARMA

CARMA implements the process description of entomological problem-solving by combining a variety of distinct reasoning paradigms. In particular, CARMA uses model-based reasoning in four different ways to assist case-based reasoning for the purpose of predicting forage loss. First, a model of grasshopper developmental periods is used to estimate probable emergence dates of nymphs and adults in order to factor the new case into distinct subcases based on feeding patterns. Second, a model of grasshopper attrition is used in temporal projection to simulate the attrition that would have occurred during the interval between the developmental distributions of the new case and the prototypical case. Third, featural adaptation constitutes a linear approximation of the function from derived case features to consumption amounts in the neighborhood of each prototypical case. Finally, critical period adjustment modifies the prediction estimate to take account of any difference in overlap between grasshopper lifespans and the critical forage growing season.

CARMA's implementation of the process model promotes the four characteristics of human expert performance mentioned above: speed, opportunism, graceful degradation, and causal explanations. CARMA is fast because, like a human expert, it can use compiled knowledge in the form of cases and rules, rather than relying entirely on computation-intensive simulations. CARMA is opportunistic in that it can recognize when no more information is required from the user (*e.g.*, when no accurate prediction can be made, or when it is too late in the season for treatment to be economical).

Graceful degradation is achieved by CARMA in two ways. First, CARMA uses multiple levels of rules ordered by certainty to infer case features. Thus, precise information can be used if available, but the absence of precise information does not cause a catastrophic fall-off in accuracy. Second, CARMA's use of CBR (*i.e.*, case-based reasoning) means that incrementally less precise information will lead to incrementally less accurate matching and adaptation, but not a catastrophic inability to provide plausible predictions and advice.

Finally, CARMA is capable of providing causal explanations, notwithstanding its use of CBR, because prototypical cases correspond to stereotypical causal scenarios. Accordingly, CARMA can apply the causal explanation associated with a prototypical case to the new case that most closely matches it. This explanation can be modified in light of the model-based reasoning used in adapting the prototypical case's prediction.[4]

Related Work

The two areas of related research most relevant to the CARMA project are the application of artificial intelligence to natural resources problem-solving, and the integration of multiple artificial intelligence techniques.

AI in Natural Resources

[4]A mechanism for providing causal explanations has not yet been implemented, but will be present in the final version of CARMA.

Two principal single-paradigm approaches to designing knowledge-based systems for natural resources management have been followed. One approach has been to use rule-based reasoning to attempt to model the process of expert human reasoning, *e.g.*, Beck, Jones, & Jones (1989) and Gupta & Suryanto (1993). An alternative approach has been to use simulation techniques to derive the expert's final answer without attempting to duplicate the expert's reasoning process (Batchelor & McClendon 1989). For example, Rodell(1978) describes a large-scale numerical simulation model for grassland ecosystems. However, there is a growing recognition that no single reasoning technique is sufficient *per se* for complex natural resources problems.

Several approaches combine rules and models, using at each stage in the reasoning process the technique that is most appropriate (Tao *et al.* 1991; Stone & Schaub 1990). This permits each reasoning technique to compensate for the weaknesses of the other. For example, a rule-based system modeled after an expert might be improved by the addition of model-based knowledge that the expert lacks (Jones, Jones, & Everett 1987; Beck, Jones, & Jones 1989). Conversely, a model-based system might become easier to use (especially for a novice) when combined with rules that interpret model results, *e.g.*, COMAX (Lemmon 1986).

CARMA represents a continuation of the trend toward incorporation of multiple reasoning paradigms to more effectively model human expert performance and to compensate for incomplete or uncertain knowledge. There is psychological evidence that much of human problem-solving uses past cases (Klein & Calderwood 1988). In particular, it is our observation that entomologists and pest managers reason with prototypical cases in rangeland grasshopper management. CARMA illustrates how case-based reasoning can be integrated with other reasoning paradigms for natural resources management.

Integration of CBR and MBR

Several previous research projects have investigated the benefits of integrating case-based reasoning with model-based reasoning (MBR). However, these projects have generally assumed the existence of a correct and complete causal model. For example, CASEY (Koton 1988) performed medical diagnosis using model-based reasoning to assist both case matching and case adaptation. However, CASEY presupposed both the existence of a complete causal theory of heart disease and complete explanations of each case in terms of that theory. Because the causal model in CARMA's domain is insufficient for accurate prediction and the causal explanations associated with cases are incomplete, the assumptions underlying CASEY's matching and adaptation strategies are inapplicable to CARMA's domain.

Rajamoney & Lee (1991) used a different approach to integrating case-based reasoning with model-based reasoning termed *prototype-based reasoning*. This approach uses a library of prototypes to decompose problems into familiar subproblems. Model-based reasoning is applied to the subproblems, a consistent composition of the subproblems is determined, and model-based reasoning is applied to determine the behavior of the resulting simplified model. As with CASEY, this approach presupposes a complete and correct (though not necessarily tractable) causal model. Similarly, Goel & Chandrasekaran's (1989) use of device models to adapt design cases presupposes that the device models are complete and correct.

Feret & Glascow (1993) describe an alterative approach under which model-based reasoning is used for "structural isolation" (*i.e.*, identification of the structural components of a device that probably give rise to the symptoms of a fault). Cases are indexed by these tentative diagnoses, which are then refined using case-based reasoning. This approach, while appropriate for diagnosis, is ill-suited for behavioral prediction in the absence of faults.

CARMA's use of model-based reasoning for case matching and adaptation represents an alternative approach to integrating CBR and MBR appropriate for domains characterized by an incomplete causal model.

Integration of CBR and RBR

Several projects have combined case-based reasoning and rule-based reasoning (RBR). PROTOS (Porter, Bareiss, & Holte 1990) uses rules to reason about the degree of equivalence between features in different cases in a technique called *knowledge-based matching*. In effect, PROTOS attempts to infer matching abstract features from nonmatching observable features. An ablation study demonstrated that the use of rules for matching contributes significantly to PROTOS's performance (Porter, Bareiss, & Holte 1990). CARMA's technique of inferring abstract features in a new case in order to establish a match with a prototypical case is similar to this approach. CARMA differs in that its prototypical cases have only abstract features, whereas PROTOS past cases are described in terms of observable features.

CABARET (Rissland & Skalak 1989) and GREBE (Branting & Porter 1991) are architectures designed to permit either rules or cases to apply to problem-solving goals. GREBE uses rules to improve case matching by inferring case facts and reformulating open-textured terms. CABARET's agenda mechanism uses heuristics to choose dynamically between rule-based reasoning and case-based reasoning. GREBE and CABARET demonstrate that integration of CBR and RBR can lead to high performance in very complex domains. CARMA is similar to these systems in that it permits both CBR and RBR to apply to high-level goals (unlike PROTOS and CASEY, which use RBR only to assist case matching and adaptation). CARMA differs from CABARET and GREBE in that the process model of rangeland pest management specifies the particular goals to which CBR and RBR apply, so in this domain (unlike the legal domains of CABARET and GREBE) it is not necessary to choose dynamically between the two techniques for each goal that arises during problem solving.

ANAPRON (Golding & Rosenbloom 1991) combines cases and rules to solve problems by using cases to represent exceptions to predictions made by the rules. However, this approach is not applicable to domains such as rangeland ecosystems, where cases and models are the main predictive components.

Conclusion

Our development of CARMA represents a strategy that we believe is generally applicable to a wide range of natural resource management issues that are characterized by high complexity and limited, diverse knowledge. These issues can be distinguished along three dimensions: tasks, perspectives, and systems. In our study, we addressed the task of grasshopper pest management, from the perspective of the agricultural producer, on the rangeland ecosystem. Other tasks in the context of this perspective and system could include advising models for weed management or grazing strategies. Alternatively, we could fix the task and system while varying the perspective. For example, the control recommendations for a grasshopper outbreak on rangeland managed by the Nature Conservancy or a wildlife refuge may be expected to differ from those provided to a rancher. Finally, we could examine grasshopper management by ranchers in a different system. For example, grasshoppers have become a very serious problem on some lands set aside as part of the Conservation Reserve Program, a program in which marginally productive croplands are being restored to grasslands.

We believe the capacity to directly address and vary the dimensions of task, perspective, and system requires multiple reasoning paradigms. A purely model-based reasoning approach does not explicitly allow for different perspectives in terms of developing and optimizing management strategies. In multiple-use settings, as found on our public forests and rangelands, an approach that has the potential to incorporate multiple tasks and perspectives (*e.g.*, by varying the cases and rules) is essential. As agriculture increasingly faces new and challenging issues of sustainability, biological diversity, natural resource conservation, and environmental quality, decision-support systems must be able to deal with extreme complexity in settings characterized by diverse forms of suboptimal qualitative and quantitative knowledge. We believe that these challenges will be most effectively addressed by adapting artificial intelligence systems to employ the reasoning techniques most appropriate for a given task in order to maximally exploit available knowledge (as human decision-makers do), rather than attempting to force the information into a particular reasoning methodology. Based on our work, there is reason to believe that this strategy will lead to advisory and decision-support systems that are more effective than single-paradigm approaches.

Acknowledgements

This research was supported in part by grants from the University of Wyoming College of Agriculture, a Faculty Grant-in-Aid from the University of Wyoming Office of Research, and by NSF Career Grant IRI-9502152.

References

Allen, T. F. H., and Hoekstra, T. W. 1992. *Toward a Unified Ecology*. New York, N.Y.: Columbia University Press.

Batchelor, W., and McClendon, R. 1989. Two approaches to knowledge engineering in an insect pest management expert system. In *Proceedings of the International Summer Meeting of the ASAE/CSAE*, volume 89-7081, 1-16.

Beck, H.; Jones, P.; and Jones, J. 1989. Soybug: An expert system for soybean insect pest management. *Agricultural Systems* 30(3):269-286.

Branting, L. K., and Porter, B. W. 1991. Rules and precedents as complementary warrants. In *Proceedings of the Ninth National Conference on Artificial Intelligence*. Anaheim: AAAI Press/MIT Press.

Chadwick, D. J., and Marsh, J. 1993. *Crop protection and sustainable agriculture*. New York: Wiley Interscience.

Edwards, C.; Lal, R.; Madden, P.; Miller, R. H.; and House, G. 1990. Sustainable agricultural systems. In Gliessman, S. R., ed., *An agroecological approach to sustainable agriculture*. Ankeny, Iowa: Soils and Water Conservation Society. 152-159.

Fedra, K. 1991. Interactive modeling of environmental impacts. In *Proceedings of Automatic Control: Triennial World Congress*, 521-527.

Feret, M. P., and Glasgow, J. I. 1993. Hybrid case-based reasoning for the diagnosis of complex devices. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 168-175. Washington, D.C.: AAAI Press/MIT Press.

Goel, A., and Chandrasekaran, B. 1989. Use of device models in adaptation of design cases. In *Proceedings of the Second DARPA Case-Based Reasoning Workshop*. Morgan Kaufmann.

Golding, A., and Rosenbloom, P. 1991. Improving rule-based systems through case-based reasoning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 22-27. Anaheim: AAAI Press/MIT Press.

Gupta, C., and Suryanto, H. 1993. A knowledge-based system for insecticide management for rice crops. *Transactions of the ASAE* 36(2):585-591.

Hewitt, G. B., and Onsager, J. A. 1983. Control of grasshoppers on rangeland in the United States: a perspective. *Journal of Range Management* 36:202-207.

Joern, A., and Gaines, S. B. 1990. Population dynamics and regulation in grasshoppers. In Chapman, R. F., and Joern, A., eds., *Biology of Grasshoppers*. New York, N.Y.: Wiley. 415-482.

Jones, J.; Jones, P.; and Everett, P. 1987. Combining expert systems and agricultural models: A case study. *Transactions of the ASAE* 30(5):1308-1314.

Klein, G. A., and Calderwood, R. 1988. How do people use analogues to make decisions? In *Proceedings of the DARPA Workshop on Case-Base Reasoning*, 209-218. Clearwater, Florida: Morgan Kaufmann.

Koton, P. 1988. *Using Experience in Learning and Problem Solving*. Ph.D. Dissertation, Massachusetts Institute of Technology. Department of Electrical Engineering and Computer Science.

Lemmon, H. 1986. An expert system for cotton crop management. *Science* 233:29-33.

Lockwood, J., and Lockwood, D. 1991. Rangeland grasshopper (orthoptera: Acrididae) population dynamics: Insights from catastrophe theory. *Environmental Entomology* 20:970-980.

Lockwood, J. A., and Lockwood, D. R. 1993. Catastrophe theory: a unified paradigm for rangeland ecosystem dynamics. *Journal of Range Management* 46:282-288.

Lockwood, J. A. 1993a. Benefits and costs of controlling rangeland grasshoppers (orthoptera: Acrididae) with exotic organisms: search for a null hypothesis and regulatory compromise. *Environmental Entomology* 22:904-914.

Lockwood, J. A. 1993b. Environmental issues involved in biological control of rangeland grasshoppers (orthoptera: Acrididae) with exotic agents. *Environmental Entomology* 22:503-518.

Lockwood, J. A. 1994. Population ecology of grasshoppers. In Gangwere, S. K., ed., *Bionomics of Grasshoppers and their Relatives*. In press.

Miller, C. 1992. Issues in sustainable agriculture: what are the next steps? *Michael Fields Agricultural Institute Bulletin 2.*

Pimm, S. L. 1991. *The Balance of Nature: Ecological Issues in the Conservation of Species and Communities.* Chicago: University of Chicago Press.

Porter, B. W.; Bareiss, E. R.; and Holte, R. C. 1990. Concept learning and heuristic classification in weak-theory domains. *Artificial Intelligence* 45(1-2).

Rajamoney, S., and Lee, H. 1991. Prototype-based reasoning: An integrated approach to solving large novel problems. In *Proceedings of the Ninth National Conference on Artificial Intelligence*. Anaheim: AAAI Press/MIT Press.

Rissland, E. L., and Skalak, D. B. 1989. Combining case-based and rule-based reasoning: A heuristic approach. In *Eleventh International Joint Conference on Artificial Intelligence*, 524-530.

Rodell, C. F. 1978. Simulation of grasshopper populations in a grassland ecosystem. In Innis, G. S., ed., *Grassland Simulation Model*. New York, N.Y.: Springer-Verlag. 127-154.

Stone, N., and Schaub, L. 1990. A hybrid expert system/simulation model for the analysis of pest management strategies. *AI Applications* 4(2):17-26.

Tao, Y.; Thompson, T.; Moser, L.; Waller, S.; Klopfenstein, T.; Ward, J.; and Wilkerson, V. 1991. Hybrid expert system for beef-forage grazing management. *Applied Engineering in Agriculture* 7(2):262-272.

	Case8	New case		Case8 after projection
		SubcaseA	SubcaseB	
Overwintering type	nymph	nymph	egg	nymph
Feeding types	grass 100%	grass 100%	grass 40% mixed 60%	grass 100%
Average period	2.0	1.2	7.0	1.2
Density	12.0	13.0	7.0	13.3
Date	September 8	August 20		September 2
Precipitation	normal	dry		normal
Temperatures	normal	cool		normal
Infest. history	mod-low	mod		mod-low
Range value	mod	high-mod		mod
Forage loss	15% (mod-low)	?		15% (mod-low)

Table 1: Case examples.

Current date	June 8
Treatment date	June 15
Treatment type	Malathion
Year-long forage loss (no treat)	40 - 60%
Year-long forage need	60%
Year-long competition (no treat)	0 - 20%
Pre-treatment forage loss	20 - 30%
Post-treatment forage loss (no treat)	20 - 30%
Treatment efficacy	80 - 90%
Post-treatment forage loss (treat)	2 - 6%
Year-long forage loss (treat)	22 - 36%
Year-long competition (treat)	0 - 0%
Forage in competition saved by treat	0 - 20%

Table 2: Example of first-year savings calculations.

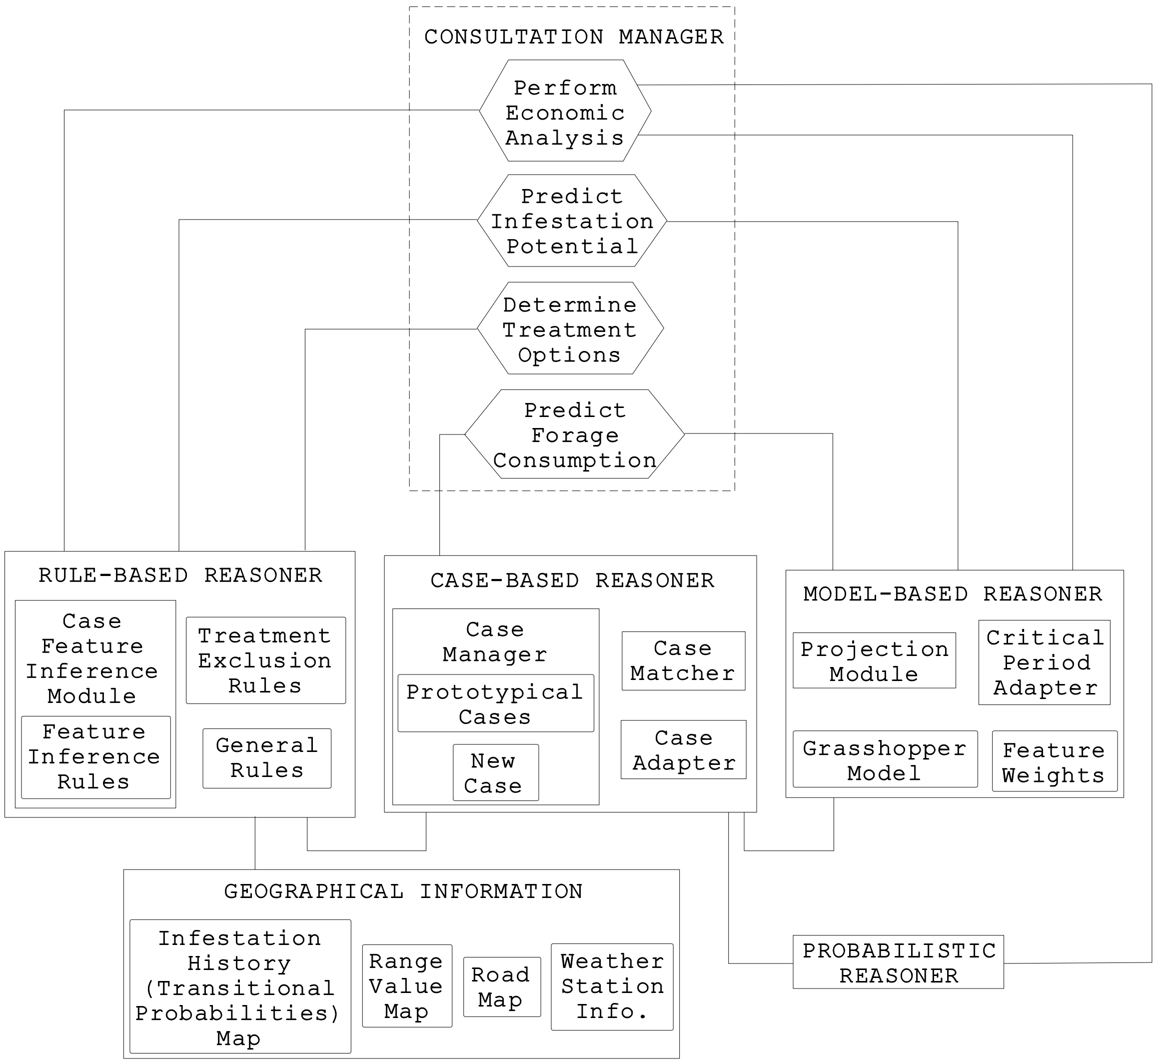

Figure 1: Organizational Overview of CARMA's Components. Hexagons represent tasks, rectangles represent modules, ovals represent information, and lines represent information paths. The order of consultation steps is not shown.

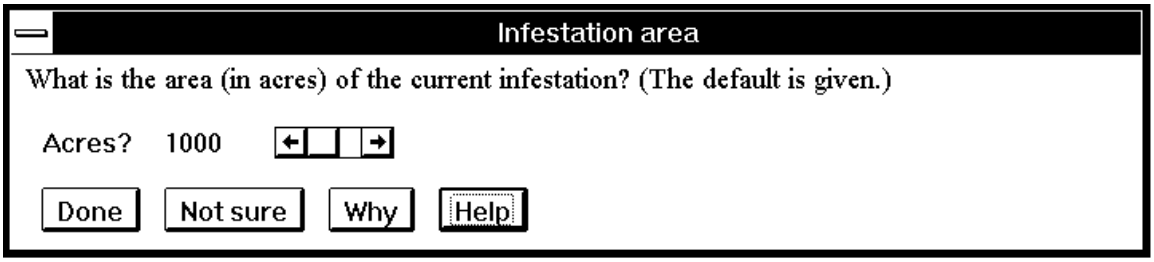

Figure 2: Interface Window for Determining the Area of Infestation.

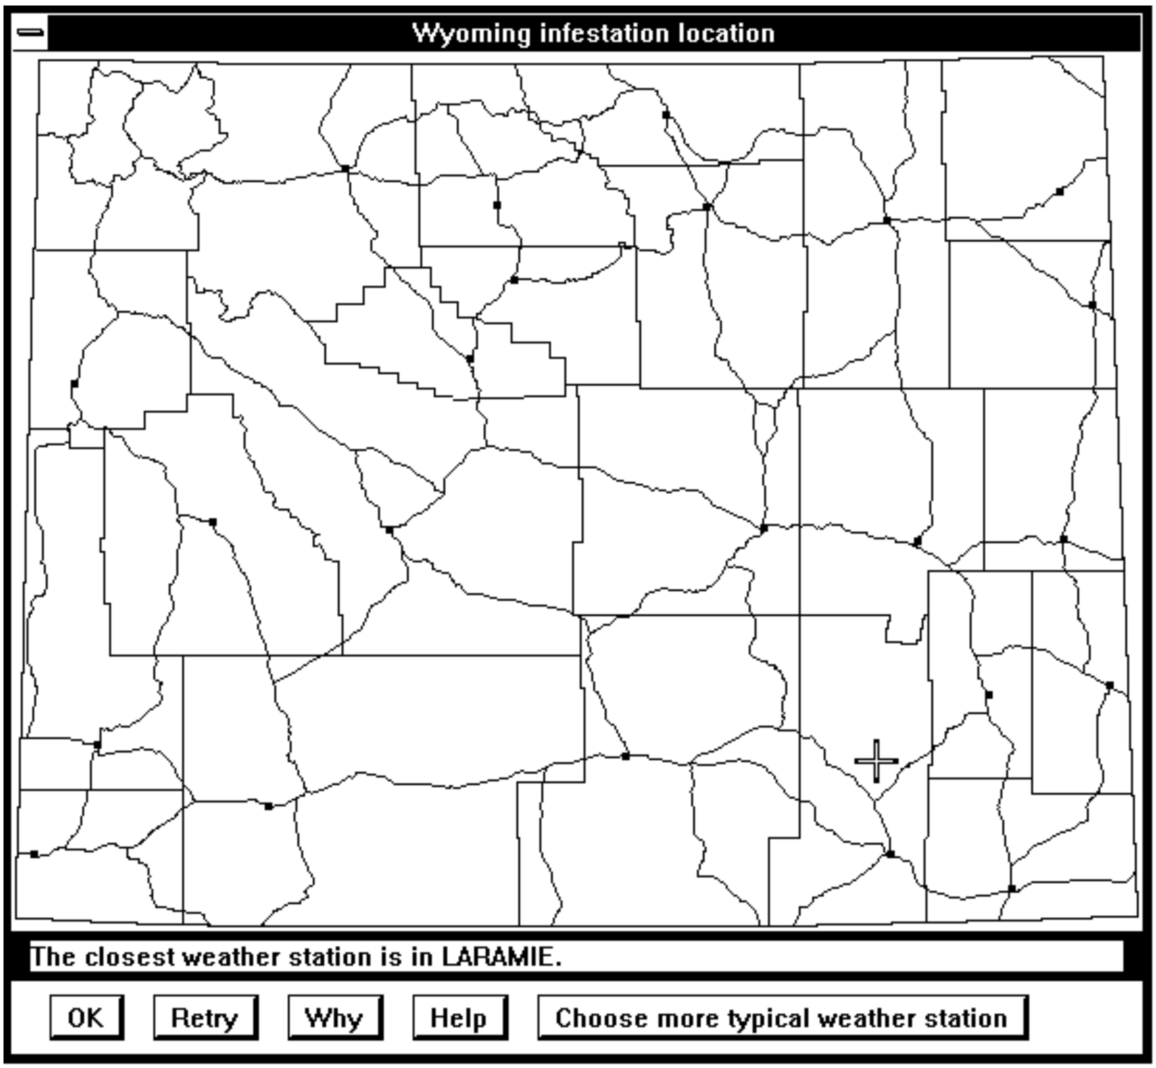

Figure 3: Interface Window for Determining Infestation Location.

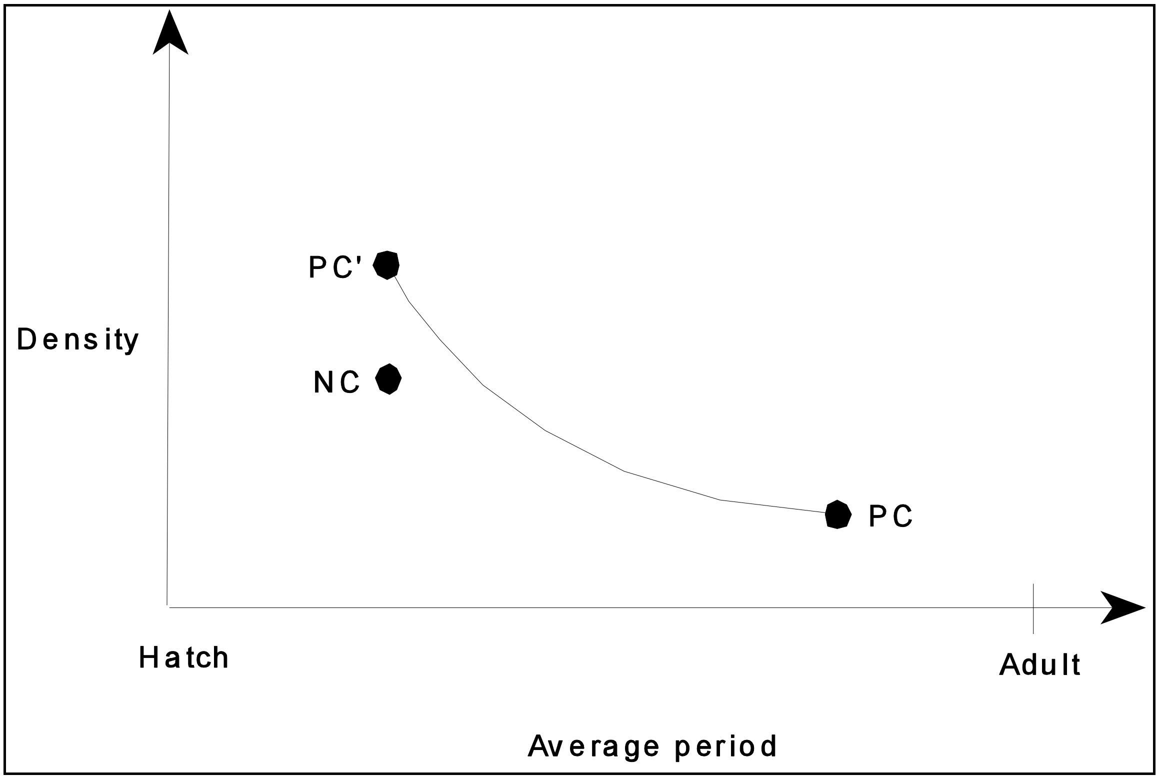

Figure 4: Projection of a prototypical case from PC to PC' to align its developmental period with new case NC.

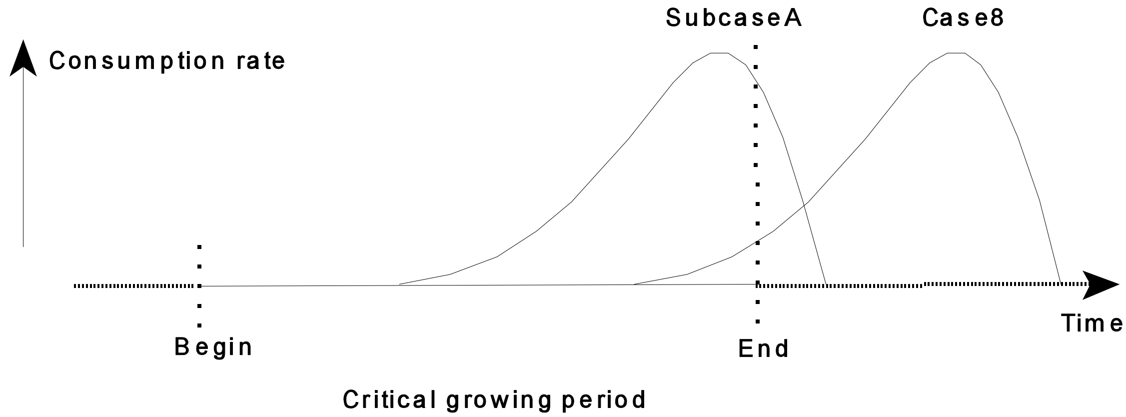

Figure 5: Critical period adjustment from Case8 to SubcaseA.

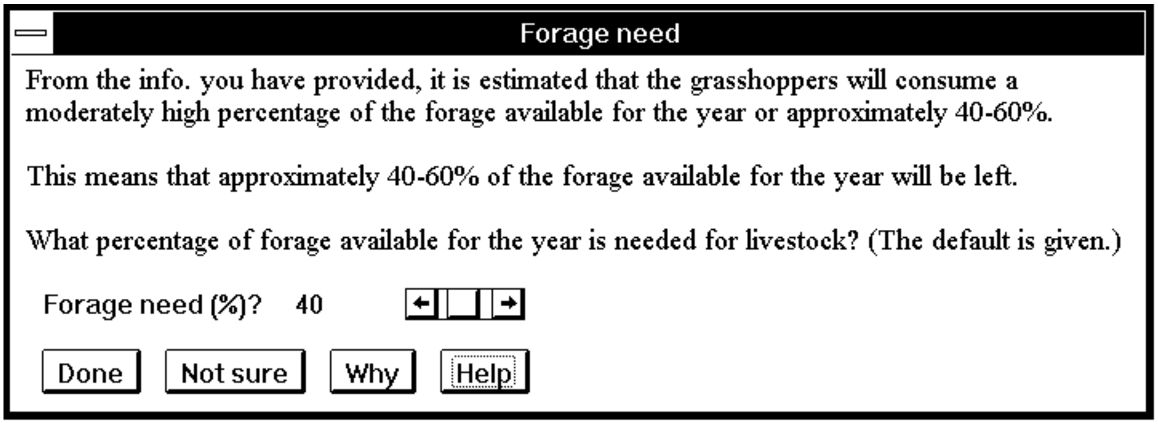

Figure 6: Interface Window for Determining Forage Need.

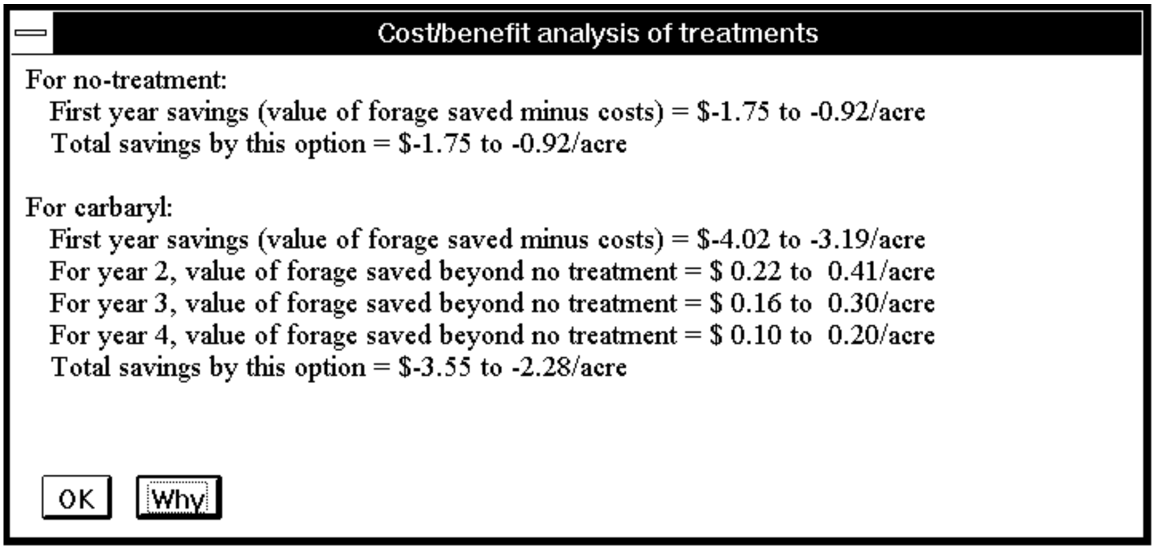

Figure 7: Cost/benefit analysis window.

Table 1: Case examples.

Table 2: Example of first-year savings calculations.

Figure 1: Organizational Overview of CARMA's Components. Hexagons represent tasks, rectangles represent modules, ovals represent information, and lines represent information paths. The order of consultation steps is not shown.

Figure 2: Interface Window for Determining the Area of Infestation.

Figure 3: Interface Window for Determining Infestation Location.

Figure 4: Projection of a prototypical case from PC to PC' to align its developmental period with new case NC.

Figure 5: Critical period adjustment from Case8 to SubcaseA.

Figure 6: Interface Window for Determining Forage Need.

Figure 7: Cost/benefit analysis window.