

CARMA

A Case-Based Rangeland Management Adviser

John Hastings, Karl Branting, Jeffrey Lockwood

■ CARMA is an advisory system for rangeland grasshopper infestations that demonstrates how AI technology can deliver expert advice to compensate for cutbacks in public services. CARMA uses two knowledge sources for the key task of predicting forage consumption by grasshoppers: (1) cases obtained by asking a group of experts to solve representative hypothetical problems and (2) a numeric model of rangeland ecosystems. These knowledge sources are integrated through the technique of model-based adaptation, in which case-based reasoning is used to find an approximate solution, and the model is used to adapt this approximate solution into a more precise solution. CARMA has been used in Wyoming counties since 1996. The combination of a simple interface, flexible control strategy, and integration of multiple knowledge sources makes CARMA accessible to inexperienced users and capable of producing advice comparable to that produced by human experts. Moreover, because CARMA embodies diverse forms of expertise, it has been used in ways that its developers did not anticipate, including pest management research, development of industry strategies, and in-state and federal pest-management policy decisions.

Grasshoppers are the most serious rangeland pest in the western United States, consuming 21 to 23 percent of rangeland forage and causing an estimated \$400 million in losses (Hewitt and Onsager 1983). Figure 1 illustrates grasshopper infestation densities in the western United States during 2000, a fairly typical year. In years of heavy infestation, grasshopper densities and economic losses might be much higher. For example, during the 1986 to 1987 outbreak, over 20 million acres of rangeland were treated for grasshoppers in the western United States at a cost of more than \$75 million.

In Wyoming, the estimated total annual loss to grasshoppers is roughly \$19 million. The southeastern quadrant of the state is particularly prone to grasshopper infestations, with significant areas of high-grasshopper densities in 30 of the last 34 years.

Various chemical and biological pesticides are available for treatment of grasshopper infestations, but the cost of using these agents often outweighs the value of the forage saved by their application. Moreover, indiscriminate pesticide application can be damaging to rangeland ecology. Despite their potential for damage, the majority of grasshopper species are usually innocuous or even beneficial to grassland ecosystems. Of more than 400 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 weed control, nutrient cycling, and food for wildlife (Lockwood 1993a, 1993b). The decision whether to use insecticides or other control measures is a complex task depending on a multiplicity of factors, including not only short-term economic costs and benefits but also preserving grasshoppers' natural enemies (Joern and Gaines 1990), safeguarding biodiversity, and protecting environmental and human health.

Before 1996, the United States Department of Agriculture (USDA) paid the entire cost of treatment on federal land, one-half the cost on state land, and one-third the cost on private land. In addition, the USDA provided intensive surveys and pest-management advice to ranchers about treatment selection. Subsequently, however, the USDA stopped providing these subsidies (except for infestations on federal rangelands that represent an immediate threat to adjacent crops) and the level of

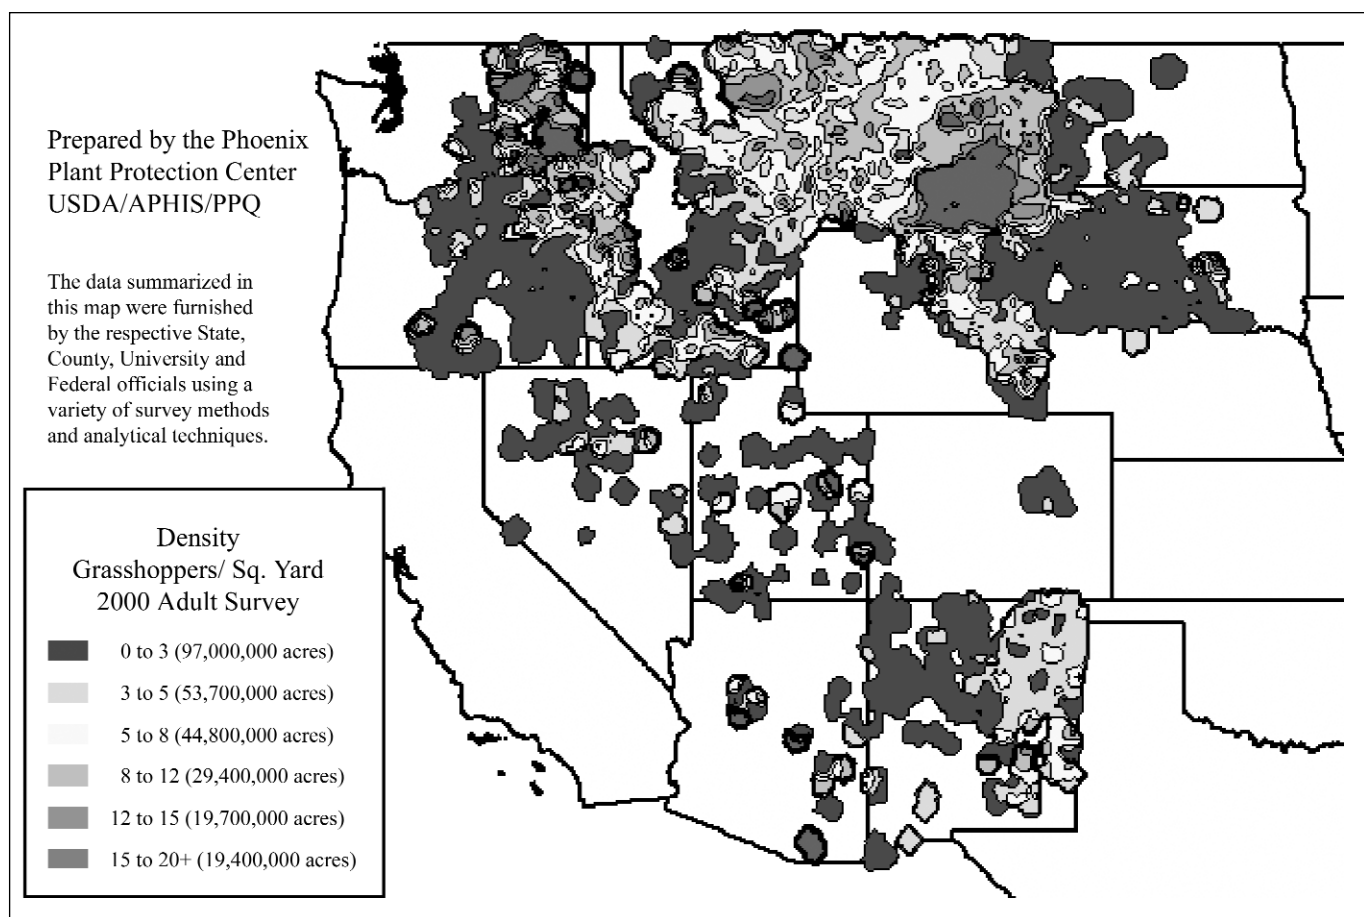


Figure 1. Western United States Rangeland Grasshopper Densities for 2000.

survey and logistical support was substantially decreased. This change in policy has increasingly shifted both the cost of treatment and the task of determining when treatment is desirable to ranchers themselves. CARMA was developed to help compensate for the decreased availability of federal assistance by helping ranchers identify and balance the factors relevant to pest-control decisions.

Task Description

CARMA's task is to help ranchers determine the most cost-effective responses to rangeland grasshopper infestations within user-defined environmental constraints. Determining the most cost-effective responses requires, at a minimum, estimating (1) the value of the forage that is likely to be consumed by grasshoppers if no action is taken, (2) the value of the portion of this forage that would be saved in current and future years under each treatment option, and (3) the cost of each option.

Estimating grasshopper forage consumption requires predicting the behavior of a rangeland

ecosystem. Various approaches to behavioral prediction are possible. In systems for which a precise model exists and accurate values of state variables can be determined, simulation can be used to predict the system's behavior. Alternatively, if there are sufficient historical data, empirical methods such as case-based reasoning (CBR) (Aamodt and Plaza 1994), decision tree induction (Quinlan 1993), or statistical techniques can be lead to accurate prediction.

Precise models exist for the behavior of many simple physical systems. However, models of agricultural, ecological, and other biological 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. Moreover, although historical data often exist for such systems, they are often insufficient for accurate prediction using empirical methods. As illustrated in figure 2, biological systems often occupy an intermediate point in the continuum between highly analytic domains, such as celestial mechanics and the prediction

of artifact behavior, and highly empirical domains, such as sociology (Allen and Hoekstra 1992). In such biological systems, both models and empirical data exist, but neither is as such sufficient for accurate prediction. Accurate prediction of the behavior of such systems requires exploitation of multiple, individually incomplete, knowledge sources.

Rangeland ecosystems typify biological systems having an extensive but incomplete causal theory and limited empirical data. Although model-based reasoning can play a role in rangeland 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 numeric simulation alone (Allen and Hoekstra 1992; Lockwood and Lockwood 1991; Pimm 1991). However, pest-management experts appear able to provide useful recommendations to ranchers, indicating that other sources of knowledge can compensate for the absence of a complete rangeland ecosystem model.

CARMA's performance objective is to emulate as closely as possible the performance of pest-management experts. The shortage of human experts makes it important for CARMA to be not only accurate but also sufficiently intuitive that it can easily be used and understood by any of its users, including ranchers, range managers (who often lack pest-management expertise), and pest managers (who might lack experience with rangeland grasshoppers).

To explicate the process whereby experts make forage-loss estimations, we performed a protocol analysis of "solve-aloud" problem solving by several experts in rangeland grasshopper management at the University of Wyoming (Hastings, Branting, and Lockwood 1996). The protocol analysis suggested that experts predict the proportion of available forage that will be consumed by grasshoppers by comparing the current situation to prototypical cases.

An example of a prototypical case is a moderate density of emerging grasshoppers in a cool, wet spring. In this situation, only a low proportion of forage is typically consumed because wet conditions both increase forage growth and promote growth of fungal pathogens that decrease grasshopper populations, and cool conditions tend to prolong the early developmental phases during which grasshoppers are most susceptible to pathogens and other mortality factors. In predicting forage consumption by comparing new cases to prototypical cases, such as the cool, wet spring prototype, experts appear to be using a form of CBR.

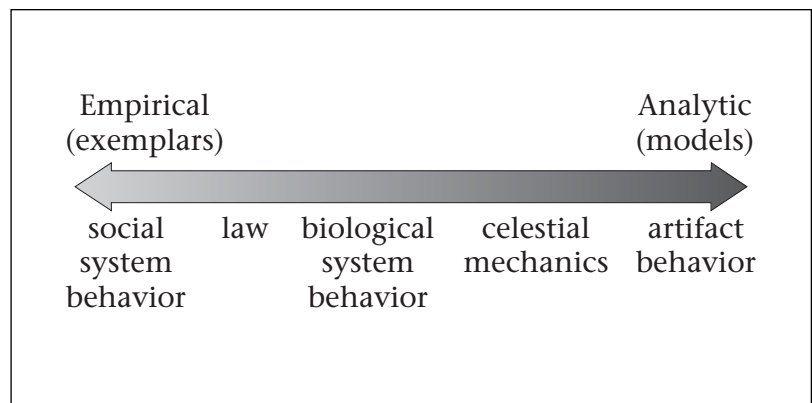


Figure 2. The Continuum from Highly Empirical Domains to Highly Analytic Domains.

If a particular new case differs in some ways from the most similar prototypical case, the expert can perform causal reasoning to adapt the prediction associated with the case to account for the differences. For example, if the population density of emerging grasshoppers in a cool, wet spring is high (rather than moderate), an expert might predict moderately low (rather than low) forage consumption because higher density generally means more consumption.

Experts seem to reason about prototypical cases in terms of abstract features that are relevant to the expert's model of rangeland ecosystems, such as grasshopper species, developmental phases, and population density. In contrast, a rancher's description is almost always in terms of directly observable features, such as the color, size, and behavior of grasshoppers; temperatures; and precipitation. As a result, determining the most similar prototypical case requires inferring the relevant abstract features from a set of observations provided by the rancher. Experts exhibit great flexibility in inferring these features. For example, if a rancher is unable to provide the information that discriminates most reliably among grasshopper species (for example, whether the grasshoppers have slanted faces or a spur on their "throats"), the expert is able to ask questions that are less reliable but easier to answer (for example, are the grasshoppers brown or green?).

If it appears that grasshoppers will consume forage needed by livestock, the expert determines which interventions are compatible with local conditions, using knowledge such as that wet conditions preclude the use of the chemical malathion and that all chemical treatments are precluded by environmental sensitivity. Finally, the expert estimates the relative value of the forage saved in this and future seasons and the cost of each control

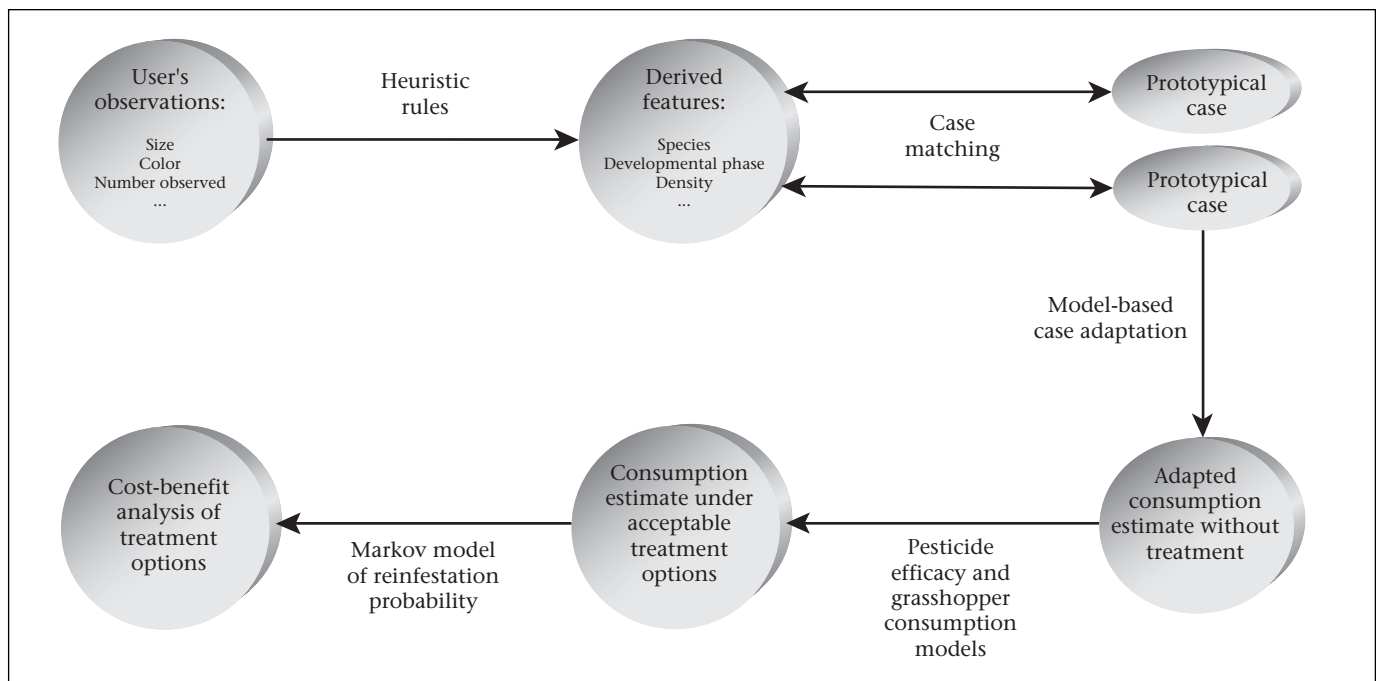


Figure 3. The Main Steps in CARMA's Consultation Process.

measure based on market price. The expert then advises the rancher to take the most economical action, either applying the most cost-effective control measure or doing nothing. Experts can justify their advice by appealing to an underlying causal model but seem to use this model only in explaining and adapting the predictions associated with prototypes and not in performing any sort of simulation.

The protocol analysis identified four important characteristics of human expert problem solving in this field: (1) graceful degradation, (2) speed, (3) causal explanations, and (4) opportunism.

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 can 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.

Speed: Human experts can provide useful advice quickly, suggesting, consistent with our process description, that human experts can use highly compiled knowledge in the form of prototypical cases.

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.

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, the protocol analysis indicated that experts in rangeland pest management use an eclectic approach that includes case-based reasoning for consumption prediction, rules for inferring case features and acceptable control measures, and causal reasoning for adaptation and explanation. Moreover, expert problem solving is fast and tolerant of inaccuracies in data.

Application Description

CARMA is designed to model the problem-solving behavior of experts in managing grasshopper infestations, as described in the previous section. CARMA emulates expert human advice by providing treatment recommendations supported by explanations in terms of causal, economic, and pragmatic factors, including a numeric estimate of the proportion of forage consumed and a cost-benefit analysis of the various treatment options.

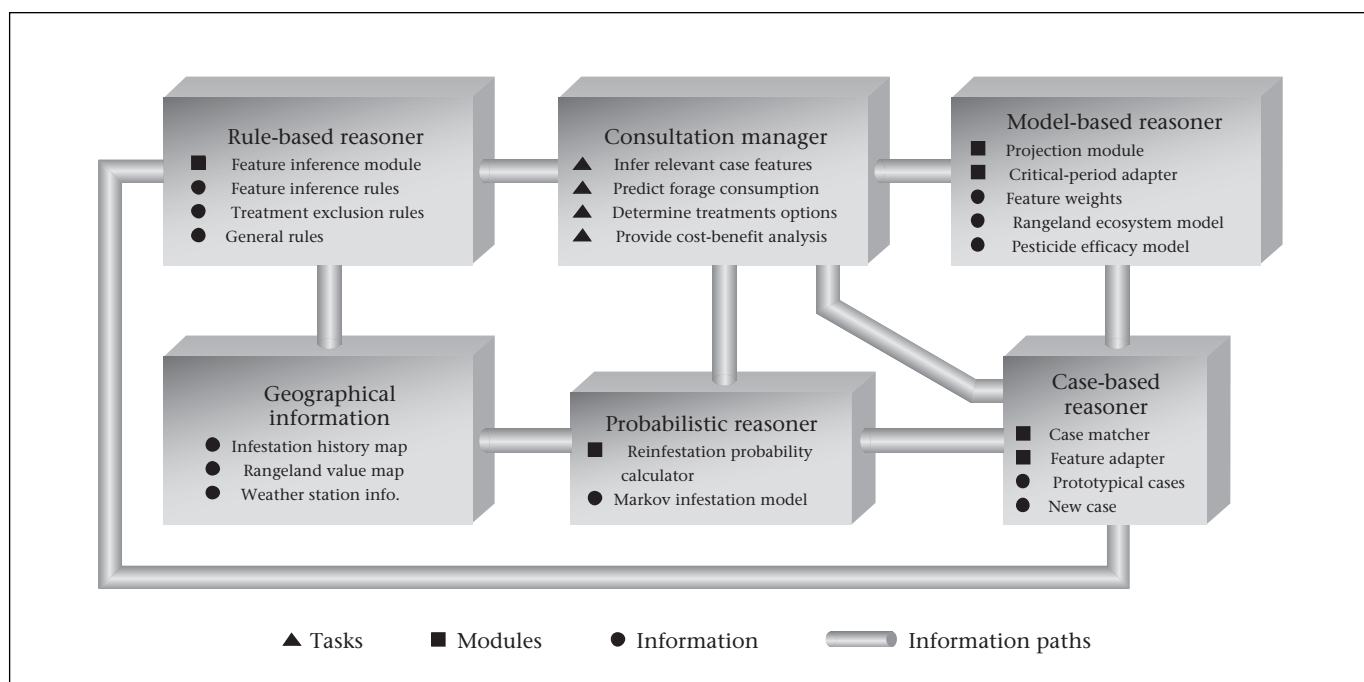


Figure 4. Organizational Overview of CARMA's Components.

Overview

CARMA's consultation process, summarized in figure 3, consists of the following steps: First, determine the relevant facts of the infestation case from information provided by the user by means of heuristic rules. Second, estimate the proportion of available forage that will be consumed by each distinct grasshopper population (that is, subcase) by matching and adapting the prototypical infestation cases that best match the facts of the current case. Third, compare total grasshopper consumption with the proportion of available forage needed by livestock. Fourth, if the predicted forage consumption will lead to economic loss, determine what possible treatment options are excluded by the case conditions. Fifth, provide an economic analysis for each viable treatment option by estimating both the first-year and long-term savings. CARMA's overall architecture is depicted in figure 4.

Determining Relevant Case Features

CARMA begins a consultation by eliciting observations from the user through a window-based interface. These observations are used to infer the relevant features of a new case, such as the species, population density, and developmental phases of the grasshoppers. CARMA uses multiple levels of rules for inferring each case feature, ordered by a qualitative estimate of each rule's accuracy or reliability. The rules are applied in succession until either the user can

provide the necessary information, or a default rule is reached.

For example, if the value of the case feature "total number of grasshoppers per square yard" is unknown to the user, CARMA instructs the user to estimate the number of grasshoppers that would be present in 18 square-foot circles (2 square yards) and divide the total by 2. If the user can't provide this information, the system attempts to infer this feature using the heuristic that grasshopper density is equal to two-thirds the number of grasshoppers seen hopping away from the user with each step taken in the field. Otherwise, the value defaults to the historic average for the area. By applying rules in order of their accuracy or reliability, CARMA reasons with the best information available.

A typical interface window for determining the observed grasshopper-type distribution appears in figure 5. It includes the options *why* for describing why this information is important to the consultation, *help* for advising the user about the various window features and their operations, *how to* to explain the proper procedure for gathering the required information, *not sure* to trigger the selection of an alternative rule for inferring the feature, *back* to return to the previous screen in the consultation, and *OK* to accept the answer chosen by the user. *Display planthopper* shows a small insect in the order *Homoptera* that resembles an immature grasshopper in both form and behavior and that the user should distinguish

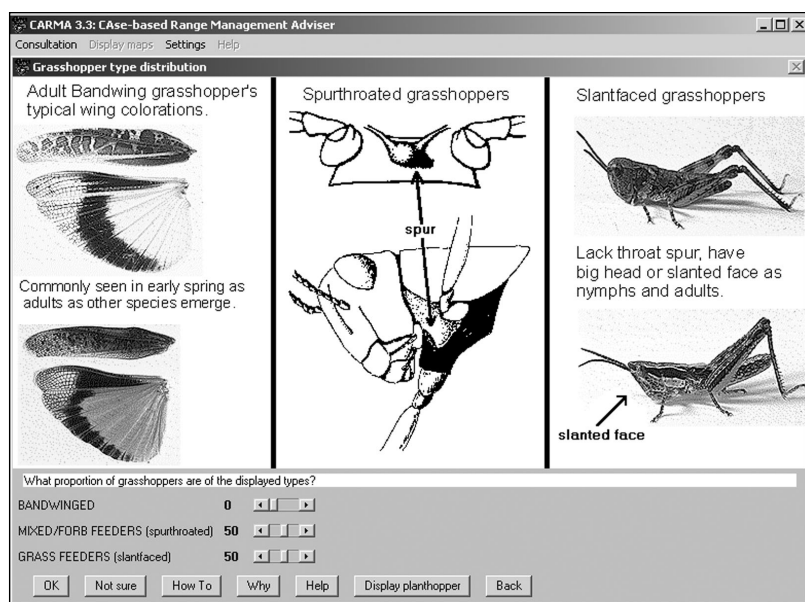


Figure 5. Elicitation of Grasshopper-Type Information in CARMA.

from a recently hatched grasshopper nymph.

Because a complete case specification is not always required for useful advice, CARMA fills in the facts of a new case opportunistically, asking 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 ended. This opportunistic policy minimizes the amount of input required for CARMA to make a decision, thereby accelerating consultations. For example, if the date and location of an infestation indicate that it is too early to assess the severity of a grasshopper infestation, CARMA advises the user to rerun the consultation at a later time without prompting for further information.

Case Matching

The protocol analysis indicated that pest managers estimate forage consumption by comparing new cases to prototypical cases. These prototypical cases differ from conventional cases in two important respects. First, the prototypical cases are not expressed in terms of observable features (for example, "Whenever I take a step, I see four grasshoppers with brightly colored wings fly"), but, rather, in terms of abstract-derived features (for example, "Approximately six nymphal-overwintering grasshoppers in the adult phase per square yard"). Second, the prototypical cases are extended in time, each representing the history of a particular grasshopper population over its life span, and stored as 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 a developmental phase in which treatment is feasible. An example prototypical case appears as case 18 in table 1.

A tract of rangeland almost invariably contains multiple grasshopper species, which can differ widely in consumption characteristics. In particular, grasshoppers that spend the winter as nymphs consume far less during the growing season than grasshoppers overwintering as eggs. CARMA therefore partitions the overall population of a new case into subcases according to overwintering type. For example, the new case set forth in table 1 is split into two subcases—(1) subcase A and (2) subcase B—based on overwintering type. Prototypical cases each represent a single grasshopper population.

To predict the forage loss of a subcase, CARMA first retrieves all prototypical cases whose overwintering type matches that of the subcase. The weighted sum of feature differences between each prototypical case and the new subcase is calculated to determine the most similar prototypical case. Match weights are determined from the mutual information gain between case features and qualitative consumption categories in a given set of training cases (Wettschereck and Dietterich 1995). Separate match weights are computed for each grasshopper overwintering type's seven case features: (1) precipitation, (2) temperature, (3) rangeland value, (4) infestation history, (5) average developmental phase, (6) density, and (7) feeding type. Quantitative features, such as density, are converted to qualitative values for computation of mutual information gain because small quantitative variations seemed to have little effect on matching. The difference between two individual feature values is determined by finding the difference between the positions of the values in an ordered qualitative feature value list. For example, rangeland value can equal one of the qualitative values in the ordered set {low, low-moderate, moderate, high-moderate, and high}, so that the matching feature difference between low and high, the maximum possible difference, is 4. The forage-loss prediction associated with the given case is then adapted to compensate for differences between the current case and the most similar prototypical case using model-based adaptation, discussed in the next subsection.

Forage-Loss Estimation

After adaptation, the consumption predictions for each subcase are summed to produce an overall consumption estimate. Because of vari-

	Case 18	New Case		Case 18 after Projection
		Subcase A	Subcase B	
Overwintering Type	Egg	Nymph	Egg	Egg
Feeding Types	Grass 90% Mixed 10%	Grass 100%	Grass 40% Mixed 60%	Grass 90% Mixed 10%
Average Phase	5.0	7.0	4.5	4.5
Density	32.0	3.0	27.0	33.3
Date	June 29	June 20		June 25
Precipitation	Normal	Dry		Normal
Temperatures	Normal	Cool		Normal
Infestation History	Moderate	High-Moderate		Moderate
Rangeland Value	Moderate-Low	Moderate		Moderate-Low
Forage Loss	90% (high)	?		90% (high)

Table 1. Case Examples.

ability resulting from the imprecise nature of rangeland ecosystems, this estimate is converted to a qualitative range (for example, high, meaning that approximately 60 to 100 percent of the available forage will be lost). The window explaining estimated forage loss, shown in figure 6, gives both aggravating and mitigating factors (that is, factors tending to increase and factors tending to reduce estimated forage loss). Explanation text in this and other CARMA windows is produced using conventional schema-based techniques (Moore 1995). 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 grasshoppers will cause economic losses.

Determining Treatment Options

If grasshoppers will cause economic losses, CARMA applies a set of rules to determine the treatment options that are excluded by the conditions of the case. Some of the information

necessary for determining exclusion is already known from the case features (for example, the presence of grasshoppers in the first nymphal instar—the earliest, readily observable developmental phase that hatches from the egg pod, which lies beneath the soil surface, into the above-ground environment—indicates an ongoing hatch, which precludes malathion and carbaryl bait from consideration). Other conditions must be determined from further user input (for example, will it be hot at the time of treatment? If so, exclude malathion.).

Treatment Recommendation

For each acceptable treatment option, CARMA provides estimates of the reduced probability of future reinfestation and current-year and long-term savings. From the estimated savings, CARMA recommends the treatment or treatments that are most economical.

CARMA calculates the total reduced probability of future reinfestation for each treatment type using a Markov model of infestation prob-

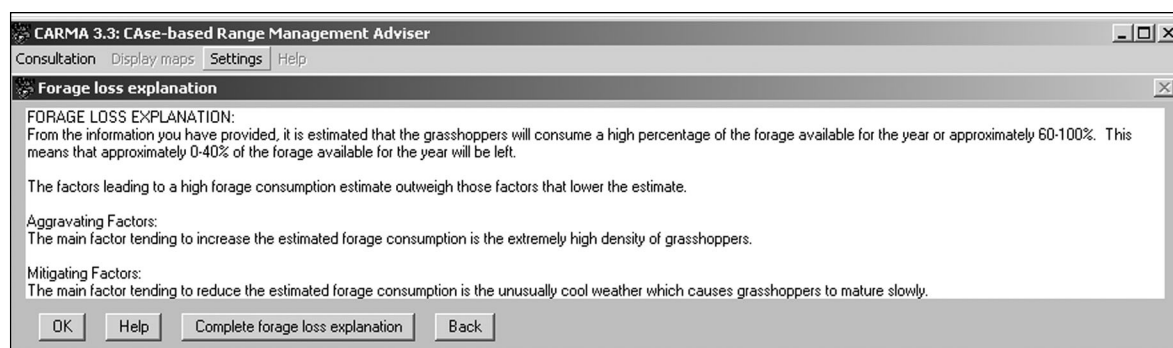


Figure 6. CARMA's Forage-Loss Estimation Window.

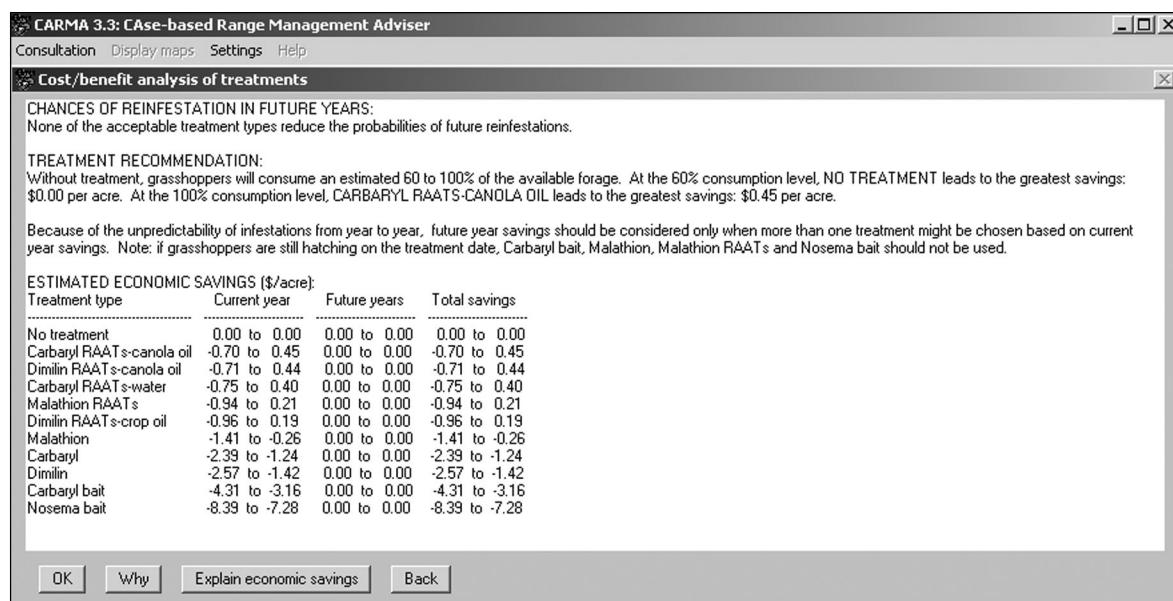


Figure 7. CARMA's Treatment Recommendation Screen.

ability for each location derived from historical data collected by the USDA and synthesized by the University of Wyoming Entomology Section (Lockwood and Kemp 1987). CARMA computes the current-year savings as the difference between the value of forage saved and the treatment cost. CARMA calculates the savings for future years for each treatment type by multiplying the reduced probabilities of reinfestation by the estimated forage loss for each subsequent year.

A typical treatment recommendation window, including estimates of future reinfestation and economic savings, appears in figure 7. CARMA lists both worst- and best-case scenarios for most calculations. Note that this analysis includes *no treatment* as an option and that negative savings indicate a loss.

CARMA recommends the treatment that is estimated to save the most under a worst-case scenario and the treatment that is estimated to

save the most under a best-case scenario. Usually, the worst- and best-case scenarios produce the same recommended treatment. Following the treatment recommendation, the initial phase of the consultation is complete.

Optionally, the user can then rerun the consultation with one or more case facts or treatment parameters altered. To facilitate the altering of treatment parameters, CARMA includes a treatment-matrix window, shown in figure 8, that permits the user to change the default values of any of the variables (for example, cost of chemical, cost of carrier, rate of application, cost of application, expected efficacy) that determine the cost of an acre protected of each treatment option or the cost of adding entirely new treatments. Although calculation of cost of an acre protected is straightforward, experience has shown that it is a common source of errors, particularly when users are attempting to determine the costs of reduced agent-area

CARMA 3.3: Case-based Range Management Adviser

Consultation Display maps Settings Help

The following table contains treatment data. Ensure the correctness of the data before continuing. Click OK to exit and accept any changes, or Restore to revert back to the initially shown settings. Click Add to add a new treatment to the bottom of the table (you may need to use the vertical scroll bar to see the new treatment row) and enter its settings, or Remove to remove a treatment highlighted by clicking in its leftmost column. Sort rows by right-clicking on a column header and left-clicking of the sort widget above the vertical scroll bar.

Treatment name	Rate of agent (fl. oz. or lbs. per acre)		Carrier	Non-aqueous carrier rate (fl. oz. or lbs. per acre)		Coverage (% of infestation)	Agent cost (\$ per gal. or lb.)	Cost of non-aqueous carrier (\$ per gal. or lb.)	Cost of application (\$ per acre)	Efficacy % (low)	Efficacy % (high)	Exclude if:	Total cost (\$ per protected acre)
No treatment	0	oz. ▼	none	0	oz. ▼	100	0.00	0.00	0.00	0	0	...	0.00
Carbaryl	16	oz. ▼	water	16	oz. ▼	100	10.80	0.00	1.25	85	95	...	2.60
Carbaryl RAATs-water	8	oz. ▼	water	8	oz. ▼	50	10.80	0.00	1.25	80	90	...	0.96
Carbaryl RAATs-canola oil	4	oz. ▼	canola oil	8	oz. ▼	50	10.80	3.80	1.25	80	90	...	0.91
Carbaryl bait	0.5	lb. ▼	wheat bran	2.0	lb. ▼	100	1.65	0.10	3.50	70	85	...	4.52
Dimilin	1	oz. ▼	crop oil (w/ twice as much water)	8	oz. ▼	100	165.00	3.80	1.25	90	95	...	2.78

Figure 8. CARMA's Treatment-Matrix Window.

treatments (RAATs), which involve applying insecticide to only a fraction of the infested area. The treatment matrix permits entirely new parameters and products to be added to CARMA as they are developed, an important feature in view of rapid recent refinements in grasshopper management methods. Easy modification of treatment parameters permits "gaming" to explore which combination of insecticide, rate, carrier, and coverage provides the best return on investment. For example, such gaming can permit ranchers and pesticide vendors to negotiate a price for a pesticide that is economical given the current infestation level and the price of substitute forage.

Uses of AI Technology

CARMA uses AI technology in two distinct ways. First, as described earlier, CARMA's control strategy emulates human experts' speed, opportunism, explanation capability, flexibility in eliciting relevant case features through a variety of alternative heuristic rules, and ability to integrate multiple knowledge sources. Second, CARMA uses model-based adaptation for the key reasoning step of predicting the amount of forage that will be consumed by grasshoppers.

Model-based adaptation consists of using CBR to find an approximate solution and model-based reasoning to adapt this approximate solution into a more precise solution. Model-

based adaptation is useful in domains in which both cases and models are available, but neither is individually sufficient for accurate prediction. Such domains are typified by chemical or biological systems with well-developed, but imperfect, models. Model-based adaptation has been applied for bioprocess recipe planning in SOPHIST (Aarts and Rousu 1996; Rousu and Aarts 1996) for selecting colorants for plastic coloring in FORMTOOL (Cheetham and Graf 1997) and in design reuse (Goel 1991).

Model-based adaptation is appropriate for CARMA's advisory task because both empirical knowledge, in the form of cases, and a grassland ecology model are available, but neither is individually sufficient for accurate prediction of forage consumption, given the information that ranchers can typically provide.

Case-Based Reasoning in CARMA

The initial impetus for using CBR for forage consumption prediction was cognitive verisimilitude. The protocol analysis suggested that human experts in this domain reason using prototypes, which is consistent with various cognitive studies that have demonstrated that examples or prototypes often play a central role in human concept structure (Klein and Calderwood 1988; Smith and Medin 1981).

During the development of CARMA, however, CBR's ability to facilitate knowledge acquisition grew in importance. Few precise records of

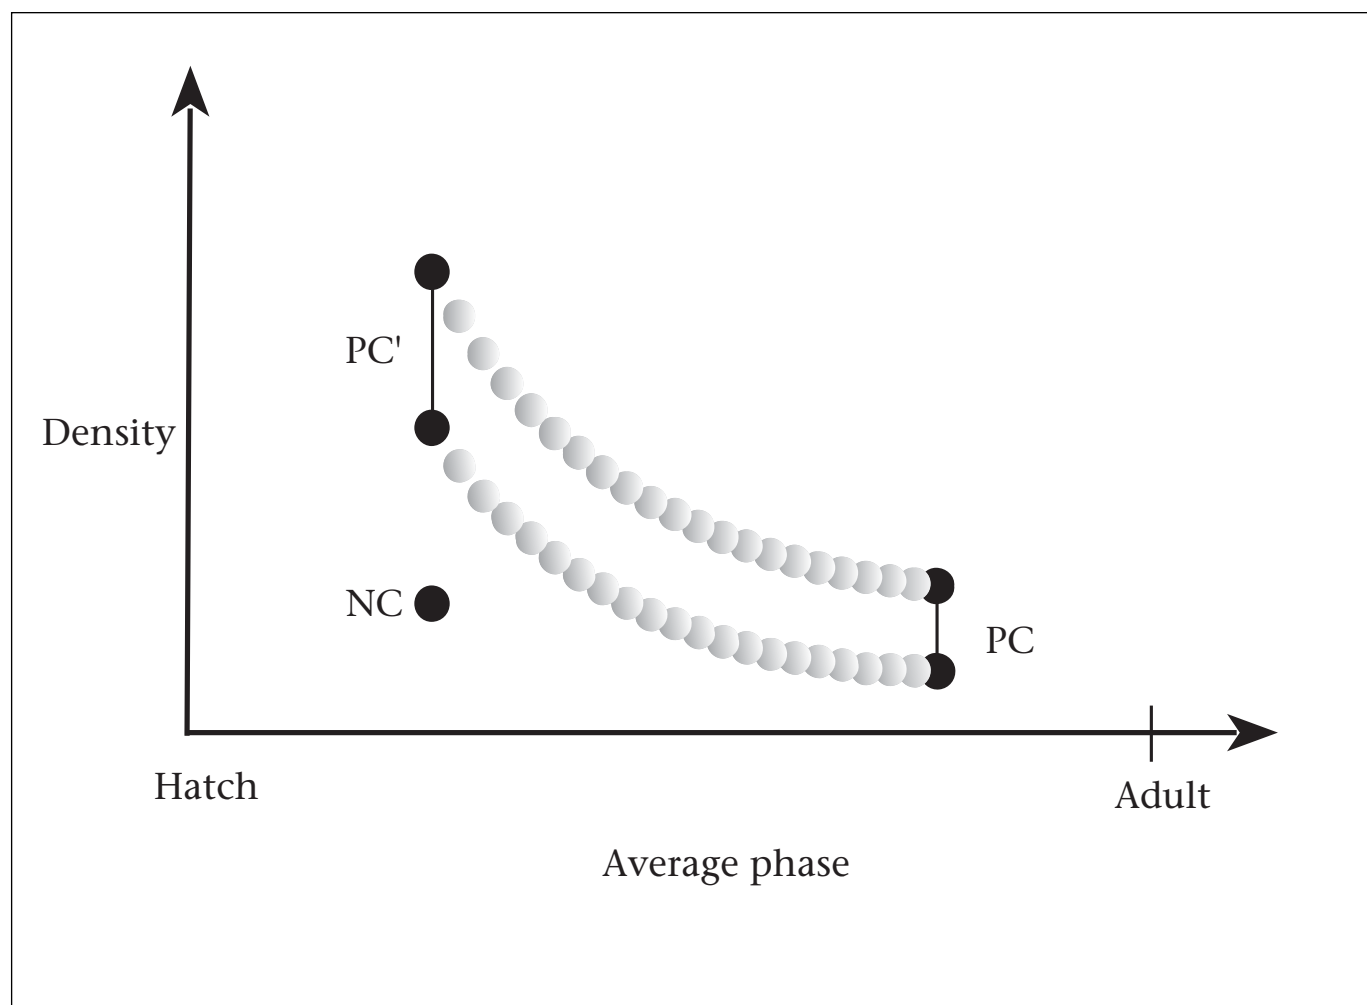


Figure 9. Projection of a Prototypical Case PC to PC' to Align Its Developmental Phase with New Case NC.

rangeland grassland infestations are available. However, there are a number of expert pest advisers with many years of experience with rangeland grasshopper infestations.

To capture human expertise, we sent questionnaires describing hypothetical Wyoming infestation cases to entomologists (including pest managers) recognized for their work in the area of grasshopper management and ecology. Each expert received 10 cases randomly drawn from a total of 20 representative cases. The descriptions of the 20 cases contained at least as much information as a rancher ordinarily provides to an entomologist. The questionnaire asked the expert to estimate the probable forage loss. Eight sets of responses were received from Wyoming experts, who had a mean of 18.0 years experience. CARMA's case library consists of the 20 hypothetical cases. The consumption prediction associated with each case is the mean of the experts' predictions for that case.

Adaptation

carma uses three techniques for adaptation: (1) temporal projection, (2) feature adaptation, and (3) critical period adaptation. Two of these techniques—(1) temporal projection and (2) critical period adaptation—make use of the rangeland ecosystem model. Temporal projection is needed because the feature values of each prototypical case are represented at a specific point in the life history of the grasshopper population. To determine the match between the grasshopper population densities of each prototypical case and a new subcase, the life history of the prototypical case must be projected forward or backward to align its average developmental phase with that of the new subcase. This projection requires a model to simulate grasshopper attrition, which depends on developmental phase, precipitation, and developmental rate (which, in turn, depends on temperature) throughout the interval of the projection.

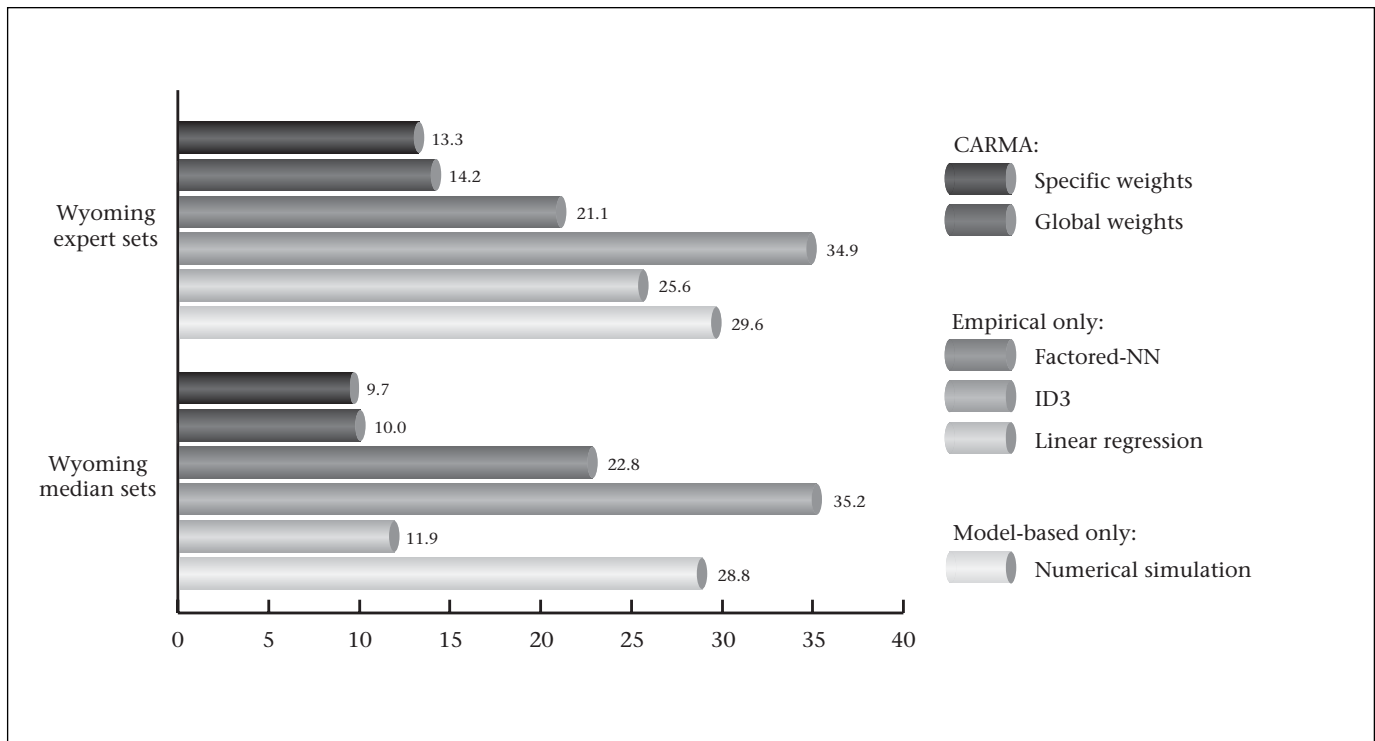


Figure 10. Root-Mean-Squared Error Rate (in Percent) for Leave-One-Out Test Results.

Figure 9 illustrates how the population in prototypical case *PC* must be projected backward in time to *PC'* to match the average developmental phase of new subcase *NC*. Projection backward in time increases grasshopper density by removing the effect of attrition over the interval of the projection, whereas projection forward in time decreases grasshopper density by adding attrition during this interval. The vertical bar corresponding to *PC* and *PC'* indicates the confidence range for grasshopper density, which always increases (indicating greater uncertainty) as a function of the interval projected.

In feature adaptation, the forage loss predicted by the best-matching prototypical case is modified to account for any feature differences (other than average developmental phase) between it and the subcase. The modification is a linear function of the feature differences. The coefficients of the linear function are determined by a form of introspective learning (Hanney and Keane 1997; Leake, Kinley, and Wilson 1995), consisting of hill climbing through parameter space to optimize leave-one-out predictive accuracy within the case library (Branting, Hasting, and Lockwood 1997).

Critical-period adaptation is needed because grasshopper consumption is most damaging if it occurs during the portion of the growing sea-

son during which forage losses cannot fully be replaced by forage growth, termed the *critical period*. The forage loss predicted by a prototypical case must be adapted if the proportion of the life span of the grasshoppers overlapping the critical period in the new case differs from that in the prototypical case. This adaptation requires determining, for both the new case and the prototypical case, the proportion of the grasshopper population's lifetime consumption occurring in the critical period. For a more complete description of model-based adaptation in *CARMA*, see Branting, Hastings, and Lockwood (1997).

Experimental Evaluation of Model-Based Adaptation

The design of *CARMA*'s forage-consumption component was based on the hypothesis that an integration of model-based and case-based reasoning can lead to more accurate forage-consumption predictions than the use of either technique individually. This hypothesis was based on the observation that neither the causal model nor the empirical data available for rangelands are individually sufficient for accurate prediction. To test this hypothesis, an ablation study was performed in which *CARMA*'s empirical and model-based knowledge components were each tested in isolation and the results compared to the performance of the full

CARMA prediction system under both global and case-specific adaptation weight modes.

Each predictive method was tested using a series of leave-one-out tests in which a set of cases (S) from a single expert was split into one test case (C) and one training set ($S - C$). The methods were trained on the forage-loss predictions of the training set and tested on the test case. This method was repeated for each case within the set (S).

CARMA's empirical component was evaluated by performing leave-one-out tests for CARMA's forage-consumption module with all model-based adaptation disabled. CARMA's forage-consumption module with model-based adaptation disabled is termed *factored nearest-neighbor prediction* (factored-NN) because under this approach, prediction is based simply on the sum of nearest-neighbor predictions for each subcase. Two other empirical methods were evaluated as well: (1) decision tree induction using ID3 (Quinlan 1986) and (2) linear regression using QR factorization (Hager 1988) to find a least squares fit to the feature values and associated predictions of the training cases.

The predictive ability of CARMA's model-based component in isolation was evaluated by developing a numeric simulation based on CARMA's model of rangeland ecology. This simulation required explicit representation of two forms of knowledge implicit in CARMA's cases: (1) the forage to an acre based on the rangeland value of the location and (2) the forage typically eaten in a day by each grasshopper for each distinct grasshopper overwintering type and developmental phase.

The accuracy of each approach was evaluated using leave-one-out testing for the responses from each of the eight Wyoming experts and for a data set consisting of the median of the predictions of the Wyoming experts on each case. The full CARMA prediction system was tested using both global-adaptation weights (CARMA-GLOBAL) and case-specific adaptation weights (CARMA-SPECIFIC).

The root-mean-squared error rates for each of the methods are set forth in figure 10. These rates provide initial confirmation for the hypothesis that integrating model-based and case-based reasoning through model-based adaptation leads to more accurate forage-consumption predictions than the use of either technique individually. The smallest root-mean-squared error rate was obtained by CARMA-SPECIFIC. On the Wyoming expert sets, the root-mean-squared error rate was 13.3 percent for CARMA-SPECIFIC and 14.2 percent for CARMA-GLOBAL. The root-mean-squared error rate was higher both for the empirical

approaches—21.1 percent for factored-NN, 34.9 percent for ID3, and 25.6 percent for linear regression—and for the purely model-based approach—29.6 percent. CARMA-SPECIFIC and CARMA-GLOBAL were also more accurate than the alternative methods on the Wyoming median set, although linear regression was only slightly less accurate.

The initial confirmation of the hypothesis that integrating model-based and case-based reasoning through model-based adaptation leads to more accurate forage-consumption predictions than the use of either technique individually is tentative because the relatively low level of agreement among experts and the absence of any external standard give rise to uncertainty about what constitutes a correct prediction. A detailed description of the empirical evaluation of CARMA is set forth in Branting, Hastings, and Lockwood (1997).

Application Use and Payoff

In June 1996, CARMA 2.0 was distributed to the University of Wyoming Cooperative Extension Offices and Weed and Pest District Offices in each of Wyoming's 23 counties and was made available for download from a University of Wyoming web site. CARMA 2.0 was used by Wyoming ranchers and pest managers every summer from 1996 to 2001. CARMA has been endorsed and advocated for use by pest managers by the United States National Grasshopper Management Board (NGMB 2001). Perhaps the greatest interest in the system has been expressed by the county-level Weed and Pest District supervisors, who—with the withdrawal of USDA support—have become the “front-line” agency in grasshopper pest management. Workshops to train these individuals in the optimal use of CARMA were developed and delivered at the request of the agency.

Although CARMA was designed as an advisory system for ranchers, CARMA's ability to robustly integrate a variety of knowledge sources led it to be applied in several ways that were not imagined when the program was developed. First, CARMA's economic analysis has been used to justify pest management policy decisions. In 1998, CARMA's economic analysis was used to generate a declaration of grasshopper disaster areas by Wyoming County Commissions, leading to low-interest, federal loans by the Farm Service Administration. CARMA's economic analysis played a role in the NGMB's recommendation of a new treatment approach, RAATs (Nelson 1999), a strategy now adopted in six states.

Second, CARMA's analysis was incorporated

into industry strategies. Uniroyal (CK Witco) developed recommendations for the use of Dimilin, a new chemical pesticide, using CARMA's analysis. Similarly, Rhône-Poulenc (Aventis) developed recommendations for the use of Fipronil based on CARMA's analysis.

Finally, CARMA-based economic analysis was incorporated into pest management research in Lockwood et al. (1999) and Lockwood and Schell (1997).

Development, Deployment, and Maintenance

CARMA was developed as a dissertation project (Hastings 1996). The out-of-pocket development costs were small, consisting of several years of graduate research assistant support and the license fees for Franz Allegro Common Lisp, the language in which CARMA was developed. However, the path to the development of CARMA was quite circuitous, with a variety of different approaches to grasshopper advising having been developed, tested, and rejected. Thus, the development costs would have been much higher outside an academic environment.

In the years since the distribution of CARMA 2.0, there have been a number of changes in pest-treatment practices. In 2001, CARMA 2.0 was updated to CARMA 3.3 to reflect these changes and include the treatment matrix discussed earlier for calculating treatment costs for each acre under various alternative economic conditions. CARMA's declarative knowledge representation made revising the program straightforward. These changes were funded by a grant from a producer of a pesticide introduced after the distribution of CARMA 2.0 and therefore not included as a treatment option in the earlier version. CARMA can be downloaded from the USDA's grasshopper-control web site or from the University of Wyoming's Grasshoppers of Wyoming and the West web site.^{1,2}

Conclusion

CARMA demonstrates how AI technology can be used to deliver expert advice to compensate for cutbacks in public services. CBR proved to be an appropriate AI technique for the forage-prediction component both because experts in this domain appear to reason with cases and because asking experts to solve example cases was an effective knowledge-acquisition technique. Model-based adaptation provided a mechanism for incorporating rangeland ecosystem models into the system without the slow performance, sensitivity to noise, and

diminished explanation capability that would have resulted from a purely simulation-based approach.

A key factor in CARMA's acceptance among users is its simple interface and speed, which make using CARMA straightforward. The combination of a simple interface, flexible control strategy, and integration of multiple knowledge sources makes CARMA accessible to inexperienced users and capable of producing advice comparable to that produced by human experts.

Notes

1. www.sidney.ars.usda.gov/grasshopper.
2. www.sdvc.uwyo.edu/grasshopper/carma.htm.

Acknowledgments

This research was supported by a grant from the University of Wyoming College of Agriculture, a faculty grant in aid from the University of Wyoming Office of Research, National Science Foundation career grant IRI-9502152, and a grant by Uniroyal (CK Witco). We gratefully acknowledge the contributions of Scott Schell to the development of CARMA's task analysis and knowledge base.

References

- Aamodt, A., and Plaza, E. 1994. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications* 7(1): 39–59.
- Aarts, R., and Rousu, J. 1996. Toward CBR for Bioprocess Planning. Paper presented at the Third European Workshop on Case-Based Reasoning (EWCBR-96), 14–16 November, Lausanne, Switzerland.
- Allen, T. F. H., and Hoekstra, T. W. 1992. *Toward a Unified Ecology*. New York: Columbia University Press.
- Branting, K.; Hastings, J.; and Lockwood, J. 1997. Integrating Cases and Models for Prediction in Biological Systems. *AI Applications* 11(1): 29–48.
- Cheetham, W., and Graf, J. 1997. Case-Based Reasoning in Color Matching. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 1–12. New York: Springer.
- Goel, A. 1991. A Model-Based Approach to Case Adaptation. Paper presented at the Thirteenth Annual Conference of the Cognitive Science Society, 7–10 August, Chicago, Illinois.
- Hager, W. 1988. *Applied Numerical Linear Algebra*. Englewood Cliffs, N.J.: Prentice Hall.
- Hanney, K., and Keane, M. 1997. The Adaptation Knowledge Bottleneck: How to Ease It by Learning from Cases. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 359–370. New York: Springer.
- Hastings, J. 1996. A Mixed-Paradigm Reasoning Approach to Problem Solving in Incomplete Causal

Theory Domains. Ph.D. thesis, Department of Computer Science, University of Wyoming.

Hastings, J.; Branting, K.; and Lockwood, J. 1996. A Multiparadigm Reasoning System for Rangeland Management. *Computers and Electronics in Agriculture* 16(1): 4–67.

Hewitt, G., and Onsager, J. 1983. Control of Grasshoppers on Rangeland in the United States: A Perspective. *Journal of Range Management* 36(2): 202–207.

Joern, A., and Gaines, S. B. 1990. Population Dynamics and Regulation in Grasshoppers. In *Biology of Grasshoppers*, eds. R. F. Chapman and A. Joern, 415–482. New York: Wiley.

Klein, G. A., and Calderwood, R. 1988. How Do People Use Analogues to Make Decisions? In *Proceedings of the DARPA Workshop on Case-Based Reasoning*, 209–218. San Francisco, Calif.: Morgan Kaufmann.

Leake, D.; Kinley, A.; and Wilson, D. 1995. Learning to Improve Case Adaptation by Introspective Reasoning and CBR. In *Case-Based Reasoning Research and Development: First International Conference on Case-Based Reasoning (ICCBR-95)*, 229–240. Lecture Notes in Artificial Intelligence 1010. New York: Springer.

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(5): 904–914.

Lockwood, J. A. 1993b. Environmental Issues Involved in Biological Control of Rangeland Grasshoppers (*Orthoptera: Acrididae*) with Exotic Agents. *Environmental Entomology* 22(3): 503–518.

Lockwood, J., and Kemp, W. 1987. Probabilities of Rangeland Grasshopper Outbreaks in Wyoming Counties. *Wyoming Agricultural Experiment Station Bulletin B(896)*.

Lockwood, J., and Lockwood, D. 1991. Rangeland Grasshopper (*Orthoptera: Acrididae*) Population Dynamics: Insights from Catastrophe Theory. *Environmental Entomology* 20(4): 970–980.

Lockwood, J., and Schell, S. 1997. Decreasing Economic and Environmental Costs through Reduced Area and Agent Insecticide Treatments (RAATs) for the Control of Rangeland Grasshoppers: Empirical Results and Their Implications for Pest Management. *Journal of Orthoptera Research* 6(1): 19–32.

Lockwood, J.; Schell, S.; Foster, R.; Reuter, C.; and Rachadi, T. 1999. Reduced Agent-Area Treatments (RAATs) for Management of Rangeland Grasshoppers: Efficacy and Economics under Operational Conditions.

International Journal of Pest Management 46(1): 29–42.

Moore, J. 1995. *Participating in Explanatory Dialogues*. Cambridge, Mass.: MIT Press.

Nelson, D., ed. 1999. Proceedings of the National Grasshopper Management Board Meeting, Colorado, 12–13 January, Denver, Colorado.

NGMB. 2001. National Grasshopper Management Board Position Statement. Paper presented at the National Grasshopper Management Board Meeting, 16–17 January, Denver, Colorado.

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


Quinlan, J. R. 1986. Induction of Decision Trees. *Machine Learning* 1(1): 81–106.

Quinlan, R. 1993. *C4.5: Programs for Machine Learning*. San Francisco, Calif.: Morgan Kaufmann.


Rousu, J., and Aarts, R. 1996. Adaptation Cost as a Criterion for Solution Evaluation. Paper presented at the Third European Workshop on Case-Based Reasoning (EWCBR-96), 14–16 November, Lausanne, Switzerland.

Smith, E. E., and Medin, D. L. 1981. *Categories and Concepts*. Cambridge, Mass.: Harvard University Press.

Wettschereck, D., and Dietterich, T. 1995. An Experimental Comparison of the Nearest-Neighbor and Nearest-Hyperrectangle Algorithms. *Machine Learning* 19(1): 5–28.



John Hastings is an assistant professor of computer science at the University of Nebraska at Kearney (UNK). He received his B.S. (1989, with honors), M.S. (1991), and Ph.D. (1996) in computer science from the University of Wyoming. After working several years as an expert systems developer and consultant, he joined the faculty at UNK in 2001. His research interests include case-based reasoning, machine learning, and applications of AI to natural resource management. His e-mail address is hastingsjd@unk.edu.



Karl Branting is a principle research scientist at LiveWire Logic, Inc., in Morrisville, North Carolina. He received a B.A. magna cum laude in philosophy from the University of Colorado, a J.D. from Georgetown University, and a Ph.D. in computer science at the University of Texas at Austin.

His current research interests include empirical natural language processing, legal applications of AI, and integration of case-based reasoning with other problem-solving paradigms. His e-mail address is karl.branting@livewirelogic.com.



Jeffrey Lockwood is a professor of entomology in the Department of Renewable Resources at the University of Wyoming. His research is concentrated on the ecology of superabundant (pest management) and rare (conservation biology) grasshopper populations. He is the executive director of The Orthopterists Society (a 350-member international organization devoted to the study of grasshoppers, crickets, and katydids) and the director of the Association for Applied Acridology International (the first and only humanitarian-based, nongovernmental organization of entomologists in the world, providing expert advice, training, and applied research to people and nations in need). Lockwood teaches courses in insect population biology, insect anatomy and physiology, and biodiversity. He has published in the areas of pest management, population modeling, biological control, agricultural ethics, and conservation biology; he serves on the editorial board of the *Journal of Insect Conservation*. His e-mail address is lockwood@uwyo.edu.