

CARMA: A Case-Based Range Management Advisor

Karl Branting
LiveWire Logic, Inc.
5500 McNeeley Drive
Suite 102
Raleigh, NC 27562
branting@livewirelogic.com

John Hastings
Dept. of Computer Science and
Information Systems
University of Nebraska at Kearney
Kearney, NE 68849
hastingsj@unk.edu

Jeffrey Lockwood
Entomology Section
Dept. of Renewable Resources
University of Wyoming
Laramie, WY 82071-3354
lockwood@uwo.edu

Abstract: CARMA is an advisory system for rangeland grasshopper infestations that demonstrates how AI technology can deliver expert advice to compensate for cut-backs in public services. CARMA uses two knowledge sources for the key task of predicting forage consumption by grasshoppers: cases obtained by asking a group of experts to solve representative hypothetical problems; and a numerical model of rangeland ecosystems. These knowledge sources are integrated through the technique of *model-based adaptation*, in which CBR 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.

Introduction

Grasshopper outbreaks cause significant economic damage to livestock producers worldwide. Grasshoppers annually consume 21–23% of rangeland forage in the western United States, causing an estimated loss of \$400 million (Hewitt and Onsager, 1983). Estimates of the value of forage lost to grasshoppers in the 17 western states in 1998 range from \$408 million (assuming replacement by leasing land) to \$1.02 billion (assuming replacement by hay) (Nelson, 1999). Similar losses were sustained in 1999. 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.

Before 1996, the USDA paid the entire cost of treatment on federal land, one-half the cost on state land, and one-third of 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 survey and logistical support was substantially decreased. CARMA was developed to help compensate for the decreased availability of federal assistance to ranchers.

Task Description

CARMA's task is to help ranchers determine the most cost-effective responses to rangeland grasshopper infestations within user-defined environmental constraints. 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 sufficiently intuitive that it can be easily used and understood by ranchers, range managers (who often lack pest management expertise), and pest managers (who may lack experience with rangeland grasshoppers). Determining the most cost-effective response to a grasshopper infestation 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.

To explicate the process whereby experts make these estimations, we performed a protocol analysis of "solve-aloud" problem solving by several experts in rangeland grasshopper management at the University of Wyoming (Hastings et al., 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 case-based reasoning

(CBR) (Aamodt and Plaza, 1994).

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 (e.g., 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 (e.g., "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 malathion and that 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 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:

- **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.
- **Speed.** Human experts can provide useful advice very quickly. This suggests that human experts can use highly compiled knowledge.
- **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.

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 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.

Application Description

Overview


CARMA's consultation process, summarized in Figure 1, consists of the following steps:

1. Determine the relevant facts of the infestation case from information provided by the user by means of heuristic rules.
2. Estimate the proportion of available forage that will be consumed by each distinct grasshopper population (i.e., subcase) by matching and adapting the prototypical infestation cases that best match the facts of the current case.
3. Compare total grasshopper consumption with the proportion of available forage needed by livestock.
4. If the predicted forage consumption will lead to economic loss, determine what possible treatment options are excluded by the case conditions.
5. Provide an economic analysis for each viable treatment option by estimating both the first-year and long-term savings.

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.

Figure 1: The main steps in CARMA's consultation process.



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). If the user can’t provide this information, the system attempts to infer this feature using the heuristic that grasshopper density is equal to 1.5 times 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 the 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 2. 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 move to the previous screen in the consultation, and “OK” to indicate that the user has chosen an answer. “Display planthopper” shows a small insect that the user should distinguish 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. 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. A tract of rangeland almost invariably contains multiple grasshopper species, which may 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. 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. 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 section.

Figure 2: Elicitation of grasshopper type information in CARMA.



Forage Loss Estimation

After adaptation, the consumption predictions for each sub-case are summed to produce an overall consumption estimate. 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 (e.g., the presence of grasshoppers in the first nymphal instar indicates an ongoing hatch, which precludes malathion and carbaryl bait from consideration). Other conditions must be determined from further user input (e.g., "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. Explanation text in this and other CARMA windows is produced using conventional schema-based techniques (Moore, 1995).

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

Figure 3: CARMA's treatment recommendation screen.



pears in Figure 3. 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 consultation is complete.

Uses of AI Technology

CARMA uses AI technology in two distinct ways. First, as described above, 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. This 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 rangeland grassland infestations are available. However, there are a number of expert pest advisors with many years of experience with rangeland grasshopper infestations.

To capture human expertise, we sent questionnaires describing hypothetical Wyoming infestation cases to entomologists 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 loss¹. Eight sets of responses were received from

¹The questionnaire also asked the expert to identify the most appropriate course of action. This information was used in the val-

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



Wyoming experts, who had a mean of 18.0 years experience². 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: temporal projection; feature adaptation; and critical period adaptation. Two of those techniques—temporal projection and 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 requires using 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 4 illustrates how the population in prototypical case PC must be projected backward in time to PC' to match the average developmental stage 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

ication of CARMA's advice, but did not become part of CARMA's knowledge base.

²The mean standard deviation of forage loss predictions for all the hypothetical cases was 12.4%, a variation sufficiently high to call into question the reliability of experts for this task. However, the issue of the validity of human expertise in this field is beyond the scope of this project. Instead, our goal was to emulate the judgments of human experts under the assumption that these judgments are valid.

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* (Leake et al., 1995; Hanney and Keane, 1997), consisting of hill climbing through parameter space to optimize leave-one-out predictive accuracy on the case library.

Critical-period adaptation is needed because grasshopper consumption is most damaging if it occurs during the portion of the growing season during which forage losses cannot be fully replaced by forage growth, termed the *critical period*. The forage loss predicted by a prototypical case must be adapted if the proportion of the lifespan of the grasshoppers overlapping the critical period in the new case differs from that in the prototypical case. This 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 et al. 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 compared the results to the performance of the full CARMA prediction system.

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: decision-tree induction using ID3³ (Quinlan, 1986) and linear regression using QR factorization (Hager, 1988) to find a least-squares fit to the feature values

³ID3 classified cases into 10 qualitative consumption categories representing the midpoints (5, 10, 15, ... , 95) of 10 equally sized qualitative ranges. ID3's error was measured by the difference between the midpoint of each predicted qualitative category and the expected quantitative consumption value.

and associated predictions of the training cases.

The predictive ability of CARMA's model-based component in isolation was evaluated by developing a numerical simulation based on CARMA's model of rangeland ecology. This simulation required explicit representation of two forms of knowledge implicit in CARMA's cases: the forage per acre based on the range value of the location, and the forage typically eaten per day per 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 for each of the methods are set forth in Table 1. These 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% for CARMA-specific and 14.2% for CARMA-global. The root-mean-squared error rate was higher both for the empirical approaches—21.1% for factored-NN, 34.9% for ID3, and 25.6% for linear regression—and for the purely model-based approach—29.6%. 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 et al., 1997).

Application Use and Payoff

In June, 1996, CARMA 2.0 was distributed to University of Wyoming Cooperative Extension Offices and Weed and Pest District Offices in each of Wyoming's 23 counties and was made available to be downloaded from a University of Wyoming website. CARMA 2.0 was used by Wyoming ranchers and pest managers every summer from 1996 to 2000. 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 National Grasshopper Management Board's recommendation of a new treatment approach, Reduced Agent Area Treatments (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, RhônePoulenc (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 of 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 2000, CARMA 2.0 was updated to CARMA 3.0 to reflect these changes and to include a spreadsheet for calculating per acre treatment costs 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. The pesticide maker's interest in being incorporated into the CARMA revision demonstrates the perception that this advisory system is an important tool in pest management.

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

Table 1: Root-mean-squared error rate (in %) for leave-one-out test results.

	CARMA		Empirical Only			Model-Based Only
	Specific weights	Global weights	Factored-NN	ID3	Linear regr.	Numerical simulation
Wyoming expert sets	13.3	14.2	21.1	34.9	25.6	29.6
Wyoming median set	9.7	10.0	22.8	35.2	11.9	28.8

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 very 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.

Acknowledgements

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, NSF Career Grant IRI-9502152, and a grant by Uniroyal (CK Witco).

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. In *Proceedings of the Third European Workshop on Case-Based Reasoning (EWCR-96)*, pages 16–27, Lausanne, Switzerland.

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*, pages 1–12, Providence, Rhode Island. Springer.

Goel, A. (1991). A model-based approach to case adaptation. In *Thirteenth Annual Conference of the Cognitive Science Society*, pages 143–148.

Hager, W. (1988). *Applied Numerical Linear Algebra*. Prentice Hall, Englewood Cliffs, New Jersey.

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*, pages 359–370, Providence, Rhode Island. Springer.

Hastings, J. (1996). *A Mixed-Paradigm Reasoning Approach to Problem Solving in Incomplete Causal Theory Domains*. PhD thesis, University of Wyoming.

Hastings, J., Branting, K., and Lockwood, J. (1996). A multi-paradigm reasoning system for rangeland management. *Computers and Electronics in Agriculture*, 16(1):47–67.

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

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*, pages 209–218, Clearwater, Florida. Morgan Kaufmann.

Leake, D., Kinley, A., and Wilson, D. (1995). Learning to improve case adaptation by introspective reasoning and CBR. In *Lecture Notes in Artificial Intelligence 1010*, pages 229–240, Sesimbra, Portugal. Springer.

Lockwood, J. and Kemp, W. (1987). Probabilities of rangeland grasshopper outbreaks in Wyoming counties. *Wyoming Agr. Exper. Sta. Bull.*, B(896). externally reviewed.

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*, 6:19–32.

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*, 46:29–42.

Moore, J. (1995). *Participating in Explanatory Dialogues*. MIT Press.

Nelson, D., editor (1999). *Proceedings of the National Grasshopper Management Board meeting*, Colorado. National Grasshopper Management Board.

NGMB (2001). National grasshopper management board position statement. In *Proceedings of the National Grasshopper Management Board meeting*, Denver, Colorado.

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

Rousu, J. and Aarts, R. (1996). Adaptation cost as a criterion for solution evaluation. In *Proceedings of the Third European Workshop on Case-Based Reasoning (EWCR-96)*, pages 354–361, Lausanne, Switzerland.

Smith, E. E. and Medin, D. L. (1981). *Categories and Concepts*. Harvard University Press.