Case Adaptation Using an Incomplete Causal Model*

John D. Hastings[1], L. Karl Branting[1], and Jeffrey A. Lockwood[2]

[1] Department of Computer Science
University of Wyoming
Laramie, Wyoming 82071-3682
{hastings, karl}@eolus.uwyo.edu

[2] Department of Plant, Soil, and Insect Sciences
University of Wyoming
Laramie, Wyoming 82071
lockwood@uwyo.edu

Abstract. This paper describes a technique for integrating case-based reasoning with model-based reasoning to predict the behavior of biological systems characterized both by incomplete models and insufficient empirical data for accurate induction. This technique is implemented in CARMA, a system for rangeland pest management advising. CARMA's ability to predict the forage consumption judgments of 15 expert entomologists was empirically compared to that of CARMA's case-based and model-based components in isolation. This evaluation confirmed the hypothesis that integrating model-based and case-based reasoning through model-based adaptation can lead to more accurate predictions than the use of either technique individually.

1 Introduction

Many types of diagnostic, monitoring, and planning tasks require prediction of the behavior of physical systems. Precise models exist for the behavior of many simple physical systems. However, models of biological, ecological, and other natural 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. Empirical methods, such as case-based reasoning, decision-tree induction, or statistical techniques, can be used for prediction if sufficient data are available. In practice, however, many biological systems are characterized both by incomplete models and insufficient empirical data for accurate induction. Accurate prediction of the behavior of such systems requires exploitation of multiple, individually incomplete, knowledge sources.

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

This paper describes the use of model-based adaptation as a technique for integrating case-based reasoning with model-based reasoning in domains in which neither technique is individually sufficient for accurate prediction. The next section describes *rangeland pest management*, a task that requires predicting the behavior of a complex biological system, and sets forth a process description of expert problem solving in this domain. Section 3 briefly describes CARMA, a system that implements this process description, and describes how CARMA performs model-based case adaptation. Section 4 describes how CARMA learns match and adaptation weights. An experimental evaluation in which the predictive accuracy of CARMA's model-based adaptation component is compared to that of case-based and model-based reasoning in isolation is set forth in Section 5. This evaluation confirms that model-based case adaptation can lead to more accurate simulation of entomologists' predictions than empirical or model-based reasoning alone.

2 Rangeland Pest Management

Rangeland ecosystems typify biological systems having an extensive but incomplete causal theory and limited empirical data. Management tasks for rangelands include optimal stocking rates and grazing systems, water development, wildlife enhancement, noxious weed control, and insect pest management. Each of these management tasks requires evaluating alternative actions by predicting their potential consequences.

The particular rangeland management task of interest to us is pest management. On average, grasshoppers annually consume 21–23% of rangeland forage in the western United States, at an estimated loss of $400 million [HO83]. Rangeland grasshopper infestations can be treated with chemical or biological insecticides, but in many situations the costs of insecticide application exceed the value of the forage saved. Determining the most cost-efficient response to a grasshopper infestation requires predicting the forage savings that would ensue from each response and comparing the savings to the cost of the response itself.

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 [LL91, Pim91, AH92]. However, entomologists and pest managers are able to provide useful recommendations to ranchers. This indicates that other sources of knowledge can compensate for the absence of a complete model of rangeland ecosystems.

Based in part on a protocol analysis of problem solving by experts in rangeland grasshopper management, we have developed the following process description of expert problem solving for this task:

1. Use rule-based reasoning to infer the relevant facts of the infestation case, such as grasshopper species, developmental phases[3], and density, from information provided by the user.

[3] During their lifetime, grasshoppers progress through three developmental stages: egg, nymph, and adult. The nymphal stage usually consists of five instars separated by

2. 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 grasshoppers using case-based and model-based reasoning.
   (b) Compare grasshopper consumption with the proportion of available forage needed by livestock.
3. If there will be competition, determine what possible treatment options should be excluded using rules such as "Wet conditions preclude the use of malathion"; "Environmental sensitivity precludes all chemical treatments."
4. If there are possible treatment options, for each one provide an economic analysis by estimating both the first-year and long-term savings using rule-based, model-based, and probabilistic reasoning.

We have implemented this problem-solving process in a system termed CARMA (CAse-based Range Management Adviser). This paper focuses on the components of CARMA that perform step 2(a), estimation of the proportion of available forage that will be consumed by grasshoppers. Making this estimation requires predicting the behavior of a rangeland ecosystem, a biological system with an incomplete causal model.

3 Model-Based Adaptation in CARMA

Our protocol analysis indicated that entomologists estimate forage consumption by comparing new cases to prototypical infestation scenarios. These prototypical cases differ from conventional cases in two important respects. First, the prototypical cases are not expressed in terms of observable features (*e.g.,* "Whenever I take a step, I see 4 grasshoppers with brightly colored wings fly"), but rather in terms of abstract derived features (*e.g.,* "Approximately 6 nymphal overwintering grasshoppers in the adult phase per square yard"). Second, the prototypical cases are extended in time, representing the history of a particular grasshopper population over its lifespan. Each prototypical case is therefore 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 a developmental phase in which treatment is feasible. An example prototypical case appears as Case8 in Table 1[4].

CARMA begins a consultation by eliciting information to infer the relevant features of a new case. When the relevant case features have been determined, CARMA can use a causal model to assist case-based reasoning in four different ways: case factoring; temporal projection; featural adaptation; and critical-period adjustment.

molts. We define the **developmental phases** of a grasshopper's lifecycle to include egg, five nymphal instars, and adult.

[4] 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.

	Case8	New case		Case8
		SubcaseA	SubcaseB	after projection
Overwintering type	nymph	nymph	egg	nymph
Feeding types	grass 100%	grass 40% mixed 60%	grass 100%	grass 100%
Average phase	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.

Factoring Cases into Subcases. CARMA's consumption prediction module first splits the overall population into subcases of grasshoppers with distinct overwintering types (*i.e.,* overwintering as nymphs or eggs), since forage consumption by those that overwinter as nymphs is much different from consumption by those that overwinter as eggs. CARMA uses a model of grasshopper developmental stages to estimate the hatch date and probable death date of each grasshopper population given the population's current developmental stage, growing season dates for the location, and current date.

Temporal Projection. To predict the forage loss of a subcase, CARMA first retrieves all prototypical cases whose overwintering types match 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 their average developmental phases with that of the new subcase. This requires using the 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. An example appears in Figure 1.

The projected prototypical case whose weighted featural difference from the new subcase is least is selected as the best match. For example, the prototypical case that best matches SubcaseA after projection is Case8, shown in Table 1. Because the developmental phase 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 phases 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 developmental phase of the grasshoppers in SubcaseA on August 20 are the same as those of Case8 two weeks later on September 2.

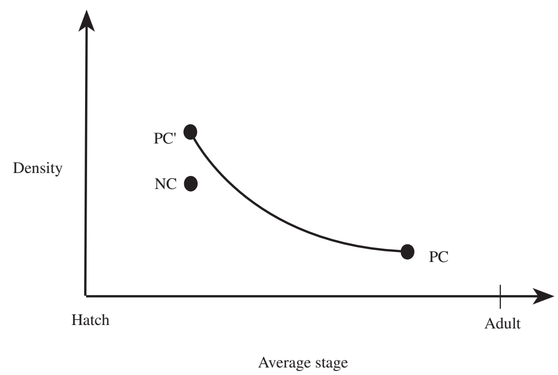

Fig. 1. Projection of a prototypical case from PC to PC′ to align its developmental phase with new case NC.

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 featural adaptation weights. For example, a lower temperature value means lower forage losses, because lower temperatures tend to slow developmental speed, 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).

Critical-Period Adjustment. 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 in the new case from the proportion in the prototypical case. This process, termed *critical-period adjustment*, requires determining the developmental phases of the new and prototypical cases that fall within the critical period and the proportion of lifetime consumption occurring in these developmental phases. The critical period of a specific parcel of rangeland is determined from the parcel's latitude and altitude.

An example of critical-period adjustment appears in Figure 2. CARMA uses a model of grasshoppers' rate of consumption at each developmental phase 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

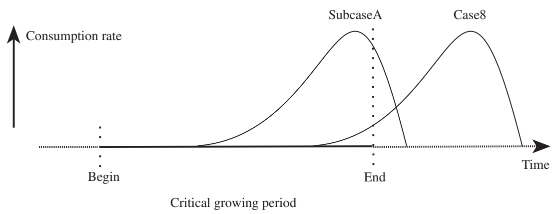

Fig. 2. Critical-period adjustment from Case8 to SubcaseA.

an overall quantitative consumption estimate. In the given case, the sum of predicted consumption of the two subcases is 57%.

In summary, CARMA can use a model of grasshopper developmental phases, consumption, and attrition, and a model of a rangeland's critical forage growth period to adapt the cases in its library. This adaptation is used both to determine the degree of relevant match between cases and to modify the consumption predictions associated with a prototypical case to apply to a new case.

4 Learning Match and Adaptation Weights

CARMA uses two sets of weights in case-based reasoning: match weights (used in the assessment of similarity between cases whose grasshopper populations have been aligned by temporal projection); and featural adaptation weights (used to adapt the consumption predicted by the best matching prototypical case in light of any featural differences between it and the subcase). General domain knowledge, such as the identifying characteristics and developmental phases of grasshoppers, can be provided by domain expert. By contrast, match and featural adaptation weights must be acquired by the system itself.

Match weights are set by determining the *mutual information gain* between case features and qualitative consumption categories in a given set of training cases, since recent research has indicated that this is often an accurate measure of featural importance for matching [WD95].

Featural adaptation weights are set by a hill-climbing algorithm, `AdaptWeights`, that incrementally varies adaptation weights A to minimize the *root-mean-squared error* (RMSE), *i.e.*,

$$\sqrt{1/n \sum_{i=1}^{n} [\texttt{PFL}(C_i, P, M, A) - \texttt{ExpertPred}(C_i)]^2}$$

for prototypical case library P and match weights M, where $\texttt{PFL}(C_i, P, M, A)$ is CARMA's predicted forage loss and $\texttt{ExpertPred}(C_i)$ is the expert's prediction of consumption for each training case C_i. CARMA can learn featural adaptation weights in either of two modes: *global*, in which a single set of weights are acquired for the entire entire case library; or *case-specific*, in which separate weights are acquired for each prototypical case.

5 Evaluation

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. To test this hypothesis, we separated CARMA's empirical and model-based knowledge components, tested each in isolation, and compared the results to the performance of the full CARMA system under both global and case-specific adaptation weight modes.

The evaluation was complicated by the absence of empirical data against which to measure CARMA's predictions. We therefore turned to expert human judgments as an external standard. To obtain a representative sample of expert opinions, we sent questionnaires to 20 entomologists recognized for their work in the area of grasshopper ecology. Each expert received 10 hypothetical cases randomly selected from a complete set of 20 cases. The descriptions of the 20 cases contained at least as much information as is typically available to an entomologist from a rancher seeking advice. The questionnaire asked the expert to make several predictions about the case, including the predicted quantitative forage loss. A total of 15 recipients of the questionnaire responded. Much to our surprise, there was a very wide variation (from 25 to 90%) in consumption predictions among the experts over the set of 20 cases. The resulting experimental case sets consist of 15 sets of expert responses containing 10 cases each.

A complication introduced by the use of expert human judgments as an evaluation standard is the possibility that in making consumption predictions human experts fail to use of all aspects of the model of grassland ecology. To test this possibility, we performed an ablation study in which we tested the effect on prediction accuracy of removing each form of adaptation knowledge from CARMA. The configuration of CARMA with the highest predictive accuracy was then compared with purely model-based and purely empirical reasoning.

5.1 Experimental Design

Each predictive method was tested using a series of leave-one-out tests in which a set of cases (S) 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). The forage loss predictions (0–100%) represent the proportion of available forage that would otherwise be available for livestock, but will instead be consumed by grasshoppers.

CARMA was tested using a protocol under which each set of training cases is used as CARMA's library of prototypical cases. This protocol is implemented in `LeaveOneOutSpecificTest` and `LeaveOneOutGlobalTest`, which perform the leave-one-out tests for the specific and global adaptation weights schemes, respectively. Both procedures call `AdaptWeights`, the hill-climbing algorithm described above. `LeaveOneOutSpecificTest` calls `AdaptWeights` with a prototypical case library containing only one case.

```
function LeaveOneOutSpecificTest(T)
    for each case C_i ϵ T do
        P := T − C_i            ;prototypical cases
        M := global match weights for set P according to info. gain
        for each prototypical case P_j ϵ P do
            T := P − P_j            ;training set
            P_j(A) := AdaptWeights(T, {P_j}, M)
        D_i := (PredictForageLoss(C_i, P, M) - ExpertPred(C_i))^2
    return ( sqrt(Avg(D)))
```

```
function LeaveOneOutGlobalTest(T)
    for each case C_i ϵ T do
        P := T − C_i            ;prototypical cases
        M := global match weights for set P according to info. gain
        G := AdaptWeights(P, P, M)
        D_i := (PredictForageLoss(C_i, P, M, G) - ExpertPred(C_i))^2
    return ( sqrt(Avg(D)))
```

Ablation Experiment. To determine the contribution of the various forms of model-based adaptation to CARMA's predictive accuracy, an ablation experiment was performed in which the performance of the full CARMA system was compared to CARMA's performance with various adaptation mechanisms disabled. The first column of Table 2 shows CARMA's average root-mean-squared error over the 15 expert sets using case specific weights (CARMA-specific). Columns two and three show CARMA-specific with, respectively, projection and critical period adjustment removed, and column four shows CARMA with featural adaptation removed. The performance of nearest-neighbor prediction (NN), *i.e.* CARMA with projection, featural adaptation, and critical period adjustment all removed,[5] is shown in column five.

These data show that full CARMA-specific actually performs worse than NN. Removing projection or featural adaptation makes performance still worse, but removing critical period adjustment makes CARMA's performance better than NN. From this, we conclude that critical period adjustment does not accurately reflect the problem-solving behavior of human experts in this predictive task.

Columns six and seven show CARMA using global weights (CARMA-global). As with CARMA-specific, CARMA-global is more accurate with critical period adjustment removed. However, CARMA-global with critical period adjustment removed, while more accurate than NN, is less accurate than CARMA-specific with critical period adjustment removed.

In summary, the ablation experiment showed that projection and featural adaptation each increased predictive accuracy but critical period adjustment decreased accuracy. Case-specific adaptation weights led to better performance

[5] Under this approach, cases are first factored into populations with distinct overwintering types, 1-NN prediction is performed for each population, and the resulting consumption predictions for all populations are summed.

than global adaptation weights. In the second experiment CARMA was therefore tested using case-specific adaptation weights and critical period adjustment disabled.

Comparison of CARMA with Empirical and Model-based Approachs. CARMA's empirical component was evaluated by performing leave-one-out-tests for a nearest-neighbor approach and two other inductive approaches that used CARMA's empirical knowledge: decision tree induction using ID3[6]; and linear approximation, which consisted of using QR factorization [Hag88] 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 numerical simulation based on CARMA's model of rangeland ecology. This simulation required 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 steps of the numerical simulation are as follows:[7]

1. Project each grasshopper population back to beginning of the growth season.
2. Simulate the density and developmental phases for each overwintering type through the end of the growth season based on the precipitation and temperature given in the case.
3. Calculate the forage eaten per day per acre based on the grasshopper density per acre and the forage eaten per day per grasshopper for each overwintering type and developmental phase as affected by temperature.
4. Convert the total forage consumed to the proportion of available forage consumed based on the forage per acre.

The effect of temperature on consumption (as a result of changing metabolism rates) was represented by multiplying a coefficient (determined from a lookup table indexed by temperature) by the forage eaten per day per grasshopper for each overwintering type. The numerical simulation was trained by hill-climbing on temperature-based coefficients to maximize the predictive accuracy on the training cases.

The accuracy of each approach was tested using leave-one-out testing for each of the 15 expert sets. The results, which appear in Table 3, include the root-mean-squared error for each of the methods.

[6] ID3 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.

[7] This model, which simulates each grasshopper population through the entire growth season, corresponds to the knowledge used by CARMA minus critical period adjustment. A simulation restricted to the critical period would correspond to the full CARMA system's knowledge.

Specific weights			No featural adaptation		Global weights	
Full	minus projection	minus CPA	minus featural adaptation	minus FA, P, CPA (NN)	Full	minus CPA
22.3	23.3	18.0	29.8	21.8	24.8	20.1

Table 2. CARMA's average percentage root-mean-squared error across 15 expert sets with various adaptation methods removed.

CARMA	Empirical Only			Model-Based Only
Specific weights minus CPA	CARMA minus FA, P, CPA (NN)	ID3	Linear appr.	Numerical simulation
18.0	21.8	29.6	31.1	25.5

Table 3. CARMA's average percentage root-mean-squared error across 15 expert sets compared with purely empirical and purely model-based approaches.

5.2 Discussion

The results of the second experiment 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 root-mean-squared error for CARMA-specific minus critical period adjustment (18.0) is 17.4% lower than for the nearest-neighbor approach (21.8) and 29.4% lower than for the numerical simulation (25.5). The error rates for the other empirical approaches on this data set were higher than for nearest-neighbor and numerical simulation: ID3 (29.6) and linear approximation (31.1). This initial confirmation is tentative because the low level of agreement among experts and absence of any external standard gives rise to uncertainty about what constitutes a correct prediction. However, this validation problem appears to be an inherent property of the domain of rangeland pest management.

Consumption prediction can be viewed as approximating a function from derived case features to consumption predictions (a *consumption function*). Prototypical cases constitute representative points in feature space for which function values are known. The prototypical cases can be used to induce a representation of the function as a decision tree (*e.g.,* ID3) or a numerical function (*e.g.,* linear approximation). The poor performance of ID3 and linear approximation suggests that the biases of these inductive methods are poorly suited to the consumption prediction task.

Alternatively, simulation can be used to derive individual values for the function. However, the incompleteness of available models of rangeland ecology limits the accuracy of this approach.

A pure nearest-neighbor approach implicitly assumes that the consumption function is constant in the neighborhood of prototypical cases. CARMA's model-based adaptation approach uses a model of rangeland ecology to attempt to approximate the consumption function in the neighborhood of individual prototypical cases. For example, projection consists of simulation through the temporal interval necessary to align the developmental phases of two cases. Although the model may be insufficient in itself for accurate consumption prediction, it may greatly improve the accuracy of nearest-neighbor prediction.

We hypothesize that CARMA-specific outperforms CARMA-global because the latter depends on the assumption that the consumption function can be approximated by a single linear equation in the neighborhood of every prototypical case. However, the poor performance of linear approximation (31.1 as compared to 21.8 for the nearest-neighbor approach) indicates that no single linear function can accurately predict consumption. Thus, it is unlikely that a single linear function is sufficient to adapt the consumption prediction of every case. CARMA-specific does not depend on the assumption that the consumption function can be approximated by a single linear equation in the neighborhood of every prototypical case. However, CARMA-specific requires a large number of cases to accurately set the adaptation weights of every prototypical case.

6 Related Work

Several previous research projects have investigated the benefits of integrating case-based reasoning with model-based reasoning. However, these projects have generally assumed the existence of a correct and complete causal model. For example, CASEY [Kot88] performed 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. Goel's use of device models to adapt design cases also presupposes that the device models are complete and correct [Goe91]. Similarly, Rajamoney and Lee's *prototype-based reasoning* [RL91] presupposes a complete and correct (though not necessarily tractable) causal model.

Feret and Glascow [FG93] 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 technique of model-based matching and adaptation represents an alternative approach to integrating CBR and MBR appropriate for domains characterized by an incomplete causal model.

7 Conclusion

This paper has described a technique for integrating case-based reasoning with model-based reasoning to predict the behavior of biological systems characterized both by incomplete models and insufficient empirical data for accurate induction. This technique is implemented in CARMA, a system for rangeland pest management advising. An empirical evaluation provided confirmation for the hypothesis that integrating model-based and case-based reasoning through model-based adaptation can lead to more accurate forage consumption predictions than the use of either technique individually.

We believe that the approach to model-based adaptation embodied in CARMA is appropriate for other domains in which empirical and model-based knowledge are each individually insufficient for accurate prediction. This approach may be particularly well suited for predictive tasks involving biological, ecological, and other natural systems.

References

[AH92] T. F. H. Allen and T. W. Hoekstra. *Toward a Unified Ecology.* Columbia University Press, New York, NY, 1992.

[FG93] M. P. Feret and J. I. Glascow. Hybrid case-based reasoning for the diagnosis of complex devices. In *Proceedings of Eleventh National Conference on Artificial Intelligence*, pages 168–175, Washington, D.C., July 11–15 1993. AAAI Press/MIT Press.

[Goe91] A. Goel. A model-based approach to case adaptation. In *Thirteenth Annual Conference of the Cognitive Science Society*, pages 143–148, 1991.

[Hag88] W. Hager. *Applied Numerical Linear Algebra.* Prentice Hall, 1988.

[HO83] G. B. Hewitt and J. A. Onsager. Control of grasshoppers on rangeland in the united states: a perspective. *Journal of Range Management*, 36:202–207, 1983.

[Kot88] P. Koton. *Using Experience in Learning and Problem Solving.* PhD thesis, Massachusetts Institute of Technology, 1988. Department of Electrical Engineering and Computer Science.

[LL91] J. Lockwood and D. Lockwood. Rangeland grasshopper (orthoptera: Acrididae) population dynamics: Insights from catastrophe theory. *Environmental Entomology*, 20:970–980, 1991.

[Pim91] S. L. Pimm. *The Balance of Nature: Ecological Issues in the Conservation of Species and Communities.* University of Chicago Press, Chicago, 1991.

[RL91] S. Rajamoney and H. Lee. Prototype-based reasoning: An integrated approach to solving large novel problems. In *Proceedings of Ninth National Conference on Artificial Intelligence*, Anaheim, July 14–19 1991. AAAI Press/MIT Press.

[WD95] D. Wettschereck and T. Dietterich. An experimental comparison of the nearest-neighbor and nearest-hyperrectangle algorithms. To appear in Machine Learning, 1995.