

Acknowledgements

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.

Appendix: Describing the contribution

A. Reasoning framework: The reasoning framework of CARMA's forage loss module is an integration of case-based reasoning with model-based reasoning through model-based adaptation. Reuse and adaptation of cases is beneficial for our application—predicting the behavior of a biological system—because, as our evaluation demonstrated, neither CBR nor MBR are as accurate individually as in combination. Our particular adaptation techniques, such as temporal projection, are extremely domain-specific. However, the general approach of using simulation to extend the coverage of cases may be widely applicable. We use a variety of forms of knowledge:

1. Prototypical cases, represented as feature vectors, expressing judgments of individual entomologists or the median of multiple judgments.
2. A model of rangeland ecology, including grasshopper development phases, consumption, reproduction, and attrition, and rangeland climate and growth patterns, obtained from an expert entomologist and from published reference materials.
3. Match weights obtained by calculation of mutual information gain between case features and qualitative forage loss categories.
4. Adaptation weights induced using a hill-climbing algorithm.

Prototypical cases are reused by adapting their forage loss predictions. Model-based adaptation requires that cases be represented using the same features as the model of rangeland ecology.

B. Task: The task of CARMA's forage loss module is determining the proportion of forage that will be consumed by grasshoppers under a given set of conditions. This is an instance of the general task of predicting the behavior of a physical system. The input to the forage loss module is a case description, and the output is a real-valued prediction between 0–100%.

C. Evaluation: Our hypotheses were: (1) an integration of model-based and case-based reasoning through model-based adaptation leads to more accurate forage consumption predictions than the use of either technique individually; and (2) the relative performance of global and case-specific adaptation weights in CARMA depends upon how closely the function from case features to forage loss can be approximated by a linear function. An empirical evaluation compared CARMA's ability to predict the forage consumption judgments of 15 expert entomologists using either global or case-specific adaptation weights to that of

CARMA's case-based and model-based components in isolation. The independent variables were case facts (cases were distributed randomly through the biologically plausible regions of instance space) and sets of expert judgments on those cases (the Expert Sets, consisting of the judgments of the 15 experts, and the Median Set, representing the median of the expert predictions). The dependent variable was the difference between the expert's and CARMA's prediction.

The primary contribution of this research is showing that an integration of CBR and MBR through model-based adaptation can increase predictive accuracy in a domain in which neither technique is individually sufficient.

References

- Branting, K., and Hastings, J. 1994. An empirical evaluation of model-based case matching and adaptation. In *Proceedings of the Workshop on Case-Based Reasoning, Twelfth National Conference Conference on Artificial Intelligence (AAAI-94)*.
- Goel, A. 1991. A model-based approach to case adaptation. In *Thirteenth Annual Conference of the Cognitive Science Society*, 143–148.
- Hager, W. 1988. *Applied Numerical Linear Algebra*. Prentice Hall.
- Hastings, J.; Branting, K.; and Lockwood, J. 1995. A multi-paradigm reasoning system for rangeland management. To appear in *Computers and Electronics in Agriculture*.
- Hewitt, G. B., and Onsager, J. A. 1983. Control of grasshoppers on rangeland in the united states: a perspective. *Journal of Range Management* 36:202–207.
- Koton, P. 1988. *Using Experience in Learning and Problem Solving*. Ph.D. Dissertation, Massachusetts Institute of Technology. Department of Electrical Engineering and Computer Science.
- Rajamoney, S., and Lee, H. 1991. Prototype-based reasoning: An integrated approach to solving large novel problems. In *Proceedings of Ninth National Conference on Artificial Intelligence*. Anaheim: AAAI Press/MIT Press.
- Wettschereck, D., and Dietterich, T. 1995. An experimental comparison of the nearest-neighbor and nearest-hyperrectangle algorithms. To appear in *Machine Learning*.

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

Table 1: 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 (-CPA)	Global weights (-CPA)	Factored-NN	ID3	Linear appr.	Numerical simulation
Expert sets	18.0	20.1	21.8	29.6	31.1	28.3
Median set	15.9	11.2	16.0	30.8	22.0	26.0

Table 2: Root-mean-squared error (in %) for leave-one-out-test results.

By contrast, learning methods without linearity bias—the numerical simulation and ID3—did not vary significantly in performance. This suggests that the Median Set is much more linear than the Expert Sets.

While case-specific adaptation weights produced better performance than global weights for the Expert Sets, the performance of specific weights on the Median Set was indistinguishable from the nearest neighbor approach, whereas global weights exhibited a significant improvement between the less linear Expert Sets to the more linear Median set, indicating that the global approach depends more on the linearity of the forage consumption function.




Figure 1: Percentage error rate improvement in Median Set over Expert Sets, $(\text{ExpertSetsErrorRate} - \text{MedianSetErrorRate}) / \text{ExpertSetsErrorRate}$, for various predictive methods.

On the Expert Sets, adaptation using global weights produced a marginal average improvement (8%) over nearest-neighbor prediction, which uses no adaptation. By contrast, adaptation using case-specific weights produced an average of 17% improvement in accuracy over nearest-neighbor prediction. Figure 2 shows the increase in accuracy of global and case-specific adapta-

tion over nearest-neighbor prediction for the 15 Expert Sets.




Figure 2: Percentage improvement of specific and global featural adaptation weight methods over nearest neighbor approach on the Expert Sets.

Conclusion

This paper has described model-based case adaptation in CARMA, a system for rangeland pest management advising. CARMA’s predictions of rangeland forage loss by grasshoppers were compared to predictions by 15 expert entomologists using either global or case-specific adaptation weights. Under both conditions, CARMA’s predictions were more accurate than CARMA’s case-based and model-based components in isolation. However, CARMA’s, case-specific adaptation weights were consistently more accurate than global adaptation weights. The experimental results suggest that case-specific adaptation weights are more appropriate in domains that are poorly approximated by a linear function.

ment 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 factored-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. Accordingly, in the evaluation of the experimental hypotheses CARMA was tested with critical period adjustment disabled.

Comparison of CARMA with Empirical and Model-based Approaches. CARMA’s empirical component was evaluated by performing leave-one-out-tests for the factored-NN approach using match weights determined from mutual information gain within each training set. This approach is equivalent to running CARMA with model-based adaptation disabled. The linearity of the Expert Sets and the Median Set was determined by linear approximation, an inductive method which consisted of using QR factorization (Hager 1988) to find a least-squares fit to the feature values and associated predictions of the training cases. An additional inductive approach to using CARMA’s empirical knowledge, decision tree induction using ID3⁵, was also tested.

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 simulation⁶ predicts the forage consumed for each day by projection based on the density, developmental phases, and forage eaten per day per grasshopper for each overwintering type, and converts that to the proportion of available forage consumed based on the forage per acre.

The steps of the numerical simulation are as follows:

⁵ID3 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.

⁶This 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.

1. Project each grasshopper population back to the beginning of critical period.
2. Simulate the density and developmental phases for each overwintering type through the end of the critical period based on the qualitative 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 metabolic 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 and the Median Set. The results, which appear in Table 2, include the root-mean-squared error for both the specific Expert Sets and the Median Set for each of the methods.

Discussion

The results of the empirical evaluation provide initial confirmation for Hypothesis 1, 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 average percentage root-mean-squared error across the 15 Expert sets for CARMA-specific minus critical period adjustment (18.0) is 17.4% lower than for the nearest-neighbor approach (21.8) and 36.4% lower than for the numerical simulation (28.3). 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.

The results of the experiment also support Hypothesis 2, that case-specific adaptation is more beneficial in less linear datasets. Figure 1 shows the difference in performance of each predictive method between the Expert Sets and the Median Set. Linear approximation had a 29% lower error rate on the Median Set (22.0 RMSE) than on the Expert Sets (31.1 RMSE).

predicted quantitative forage loss. A total of 15 recipients of the questionnaire responded. The resulting experimental case sets comprised the 15 sets of expert responses consisting of 10 cases each (the *Expert Sets*), a set of 20 cases representing the median of the experts’ prediction on each case (the *Median Set*), and ProtoL, the library of 19 prototypical cases generated by an independent entomologist.

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 a preliminary ablation study in which we tested the effect on predictive accuracy of removing each form of adaptation knowledge from CARMA. The configurations of CARMA (for both specific and global weights schemes) with the highest predictive accuracy were then used as the full CARMA system.

To test Hypothesis 1, we separated CARMA’s empirical and model-based knowledge components, tested each in isolation on the Expert and Median sets, and compared the results to the performance of CARMA using both global and case-specific adaptation weights. To test the Hypothesis 2, we determined the extent to which the Expert and Median sets could be approximated by a linear equation and compared this to the degree to which global and case-specific adaptation improved accuracy beyond simple nearest-neighbor prediction.

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 (between 0% and 100%) represent the proportion of available forage that would otherwise be available for livestock, but will instead be consumed by grasshoppers.

CARMA’s global and case-specific adaptation weight methods were initially tested using leave-one-out testing with ProtoL as its case library and the Median Set as the test set. However, CARMA’s performance was disappointing. Analysis of the source of errors revealed inconsistencies between prototypical cases in ProtoL and in the Median Set. Cases in the two sets with nearly identical features often had very different predictions. These inconsistencies could not be overcome by any set of adaptation weights and led to instability in adaptation weights: adaptation weights sometimes varied by as much as a factor of 35,000 from one leave-one-out-test to another within a single training set. We therefore abandoned ProtoL and shifted to 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)
1   for each case  $C_i \in T$  do
2      $P := T - C_i$  ;prototypical cases
3      $M :=$  global match weights for set  $P$ 
        according to info. gain
4     for each prototypical case  $P_j \in P$  do
5        $T := P - P_j$  ;training set
6        $P_j(A) :=$  AdaptWeights( $T, \{P_j\}, M$ )
7        $D_i :=$  (PredictForageLoss( $C_i, P, M$ )
        - ExpertPred( $C_i$ ))2
8   return ( $\sqrt{\text{Avg}(D)}$ )

```

```

function LeaveOneOutGlobalTest(T)
1   for each case  $C_i \in T$  do
2      $P := T - C_i$  ;prototypical cases
3      $M :=$  global match weights for set  $P$ 
        according to info. gain
4      $G :=$  AdaptWeights( $P, P, M$ )
5      $D_i :=$  (PredictForageLoss( $C_i, P, M, G$ )
        - ExpertPred( $C_i$ ))2
6   return ( $\sqrt{\text{Avg}(D)}$ )

```

Ablation Experiment. An ablation study was performed to determine the combination of adaptation knowledge sources leading to the highest predictive accuracy. In this study, the full CARMA system was compared to CARMA’s performance with various adaptation mechanisms disabled. The first column of Table 1 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 *factored nearest-neighbor prediction* (factored-NN), *i.e.*, CARMA with projection, featural adaptation, and critical period adjustment all removed,⁴ is shown in column five.

These data show that full CARMA-specific actually performs worse than factored-NN. Removing projection or featural adaptation makes performance still worse, but removing critical period adjustment makes CARMA’s performance better than factored-NN. From this, we conclude that critical period adjust-

⁴Under 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.

By contrast, match and featural adaptation weights must be acquired by the system itself.

Match Weights

Match weights are set by determining the *mutual information gain* between case features and qualitative consumption categories in a given set of training cases. Mutual information gain is used because recent research has indicated that it is often an accurate measure of featural importance for matching (Wettschereck & Dietterich 1995). Separate match weights are computed for each grasshopper overwintering type for the case features **precipitation, temperature, range value, infestation history, average developmental phase, density, and feeding type**.³

Quantitative features, such as **density**, are converted to qualitative values for computation of mutual information gain, since small quantitative variations appeared to make little contribution to matching in CARMA. The matching feature 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. The similarity of two cases is determined by summing each individual feature difference multiplied by the corresponding match weight.

Adaptation Weights

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.*,

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

for prototypical case library P and match weights M , where $\text{PFL}(C_i, P, M, A)$ is CARMA's predicted forage loss and $\text{ExpertPred}(C_i)$ is expert's prediction of consumption for each training case C_i .

Separate adaptation weights are computed for each grasshopper overwintering type for the six features **precipitation, temperature, range value, infestation history, density, and feeding type**. CARMA can learn featural adaptation weights in either of two modes: *global*, in which a single set of weights are acquired for the entire case library; or *case-specific*, in which separate weights are acquired for each prototypical case

³A match weight is necessary for developmental phase even though temporal projection typically aligns the developmental phases of the cases because temporal projection is limited to 2 weeks in the match phase. The purpose of this limitation is to bias matches towards cases with similar developmental phases. However, in the adaptation stage, any remaining misalignment in developmental stages is eliminated by continuing temporal projection until the developmental phases are completely aligned.

In computing the featural adaptation weights, qualitative case features (such as precipitation = "Dry") are converted into quantitative values based on the position of the value in an ordered qualitative feature value list. An adaptation feature difference is computed as the difference between the quantitative feature values of the two cases. The consumption prediction of the matching prototypical case is adjusted by the sum of the adaptation feature differences multiplied by the adaptation weights for each feature.

Evaluation

The design of CARMA's forage consumption component was based on the following hypothesis:

Hypothesis 1. 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 is based on the observation that neither the causal model nor the empirical data available for rangelands is individually sufficient for accurate prediction.

The task of consumption prediction can be viewed as simulating a function from case features to forage consumption. Our second hypothesis was the following:

Hypothesis 2. The relative performance of global and case-specific adaptation weights in CARMA depends upon how closely the function from case features to forage consumption can be approximated by a linear function.

The rationale for this hypothesis is that featural adaptation consists, in effect, of a linear approximation of the forage consumption function in the vicinity of each prototypical case after it has been adapted by case factoring and temporal projection. Use of global adaptation weights rests on the assumption that the forage consumption function can be approximated by the same linear function in the vicinity of every case. Use of case-specific weights, by contrast, is based on the assumption that a separate linear function is required to approximate the forage consumption function in the vicinity of each case. The danger of the latter assumption is that if there are few training cases the adaptation weights may suffer from overfitting.

Little empirical data was available to serve as a measure of 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 as much or more information than 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

of rangeland forage at an estimated loss of \$400 million (Hewitt & Onsager 1983). 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.

Rangeland pest management requires predicting the forage savings that would ensue from each response and comparing the savings to the cost of the response itself. A central step in determining the most cost-efficient response is estimating the proportion of available forage that will be consumed by grasshoppers. A protocol analysis of entomologists experienced in pest management indicated that they estimate forage consumption by comparing new cases to prototypical infestation scenarios.

CARMA's case library consists of prototypical cases that 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.

CARMA begins a consultation by eliciting information that can be used to infer the relevant features of a new case. When the relevant case features have been determined, CARMA predicts forage loss by using a causal model to assist case-based reasoning in four different ways:² case factoring; temporal projection; featural adaptation; and critical-period adjustment.

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 the forage consumption by those that overwinter as nymphs is much different from 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 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 the average developmental phase 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.

The projected prototypical case having the lowest weighted featural difference from the new case is selected as the best match.

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.

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. CARMA uses a model of grasshopper's rate of consumption at each developmental phase to calculate the proportion of lifetime consumption occurring before the end of the critical period. After adaptation, the consumption predictions for each subcase (*i.e.*, distinct overwintering types) are summed to produce an overall quantitative consumption estimate.

In summary, CARMA uses a model of grasshopper developmental phases, consumption, and attrition, and a model of a rangeland's critical forage growth period for adaptation of 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.

Learning Match and Adaptation Weights

CARMA uses two sets of weights in case-based reasoning: match weights (used to assess case similarity); 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 the domain expert.

²A detailed example of the application of these model-based adaptation techniques to a specific case appears in (Hastings, Branting, & Lockwood 1995).

Global and Case-Specific Model-based Adaptation

John D. Hastings and L. Karl Branting

Department of Computer Science
University of Wyoming
Laramie, Wyoming 82071-3682

Abstract

CARMA (CAse-based Range Management Adviser) is a system that integrates case-based reasoning with model-based reasoning for rangeland pest management. CARMA's predictions of rangeland forage loss by grasshoppers were compared to predictions by 15 expert entomologists using either global or case-specific adaptation weights. Under both conditions, CARMA's predictions were more accurate than CARMA's case-based and model-based components in isolation. However, CARMA's case-specific adaptation weights were consistently more accurate than global adaptation weights. The experimental results suggest that case-specific adaptation weights are more appropriate in domains that are poorly approximated by a linear function.

Introduction

Model-based case adaptation is a technique in which a domain model is used to assist in adapting solutions associated with past cases to apply to new problems (Goel 1991). Most previous implementations of model-based adaptation have assumed the existence of a complete and accurate domain model (Koton 1988; Goel 1991; Rajamoney & Lee 1991). The purpose of model-based case adaptation in these systems is to reduce the search associated with a complex model.

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. Moreover, such systems often lack sufficient data for effective use of empirical methods, such as case-based reasoning, decision-tree induction, or statistical techniques. In such systems, model-based case adaptation provides a mechanism to combine multiple, individually incomplete, knowledge sources to increase predictive accuracy.

This paper describes model-based case adaptation in CARMA (CAse-based Range Management Adviser), a system that integrates case-based reasoning with model-based reasoning for rangeland pest management. Rangeland pest management is a task requiring

predictions in a biological system characterized both by an incomplete model and insufficient empirical data for accurate use of empirical techniques. Part of this task requires predicting the probable forage consumption by grasshoppers in an infestation.¹

CARMA contains a detailed model of grasshopper development phases, consumption, reproduction, and attrition. However, this model contains parameters, termed "featural adaptation weights", representing the relative importance of case features to total forage consumption, whose precise values are not known *a priori*, but must be acquired inductively. Featural adaptation weights may either be *global*, applying to the entire case library, or *case-specific*, with separate weights applying to each case. Whether global or case-specific is better depends on whether the feature weights apply to the entire instance space or are instead specific to portions of feature space corresponding to one or more cases. If the former, case-specific weights might lead to overfitting. If the latter, no effective set of global weights will exist.

The next section describes how CARMA performs model-based case adaptation. CARMA's procedure for learning match and adaptation weights is then described. The final section describes an experimental evaluation in which CARMA's predictions of rangeland forage consumption by grasshoppers were compared to predictions by 15 expert entomologists using either global or case-specific adaptation weights. Under both conditions, CARMA's predictions were more accurate than CARMA's case-based and model-based components in isolation. However, case-specific adaptation weights were consistently more accurate than global adaptation weights.

Model-Based Adaptation in CARMA

Rangeland pest management is the task of determining the most cost-efficient response to a grasshopper infestation. This task is important because in the western United States grasshoppers annually consume 21–23%

¹A detailed description of the rangeland pest management task is set forth in (Branting & Hastings 1994) and (Hastings, Branting, & Lockwood 1995).