Compositional Instance-Based Learning

Karl Branting and **Patrick Broos**
Department of Computer Science, University of Wyoming
Laramie, Wyoming 82071-3682
{karl,patb}@eolus.uwyo.edu

Abstract

This paper proposes a new algorithm for acquisition of preference predicates by a learning apprentice, termed *Compositional Instance-Based Learning* (CIBL), that permits multiple instances of a preference predicate to be composed, directly exploiting the transitivity of preference predicates. In an empirical evaluation, CIBL was consistently more accurate than a 1-NN instance-based learning strategy unable to compose instances. The relative performance of CIBL and decision tree induction was found to depend upon (1) the complexity of the preference predicate being acquired and (2) the dimensionality of the feature space.

Introduction

A central impediment to the construction of knowledge-based systems is the high cost of knowledge base development and maintenance. One approach to reducing these costs is to design systems that can acquire knowledge by observing human problem-solving steps during normal use of the system. Systems that engage in this form of learning are termed *learning apprentice systems* (Mitchell, Mahadevan, & Steinberg 1985). Learning apprentice systems have been developed for VLSI design (Mahadevan *et al.* 1993), acquisition of "interface agents" (Maes & Kozierok 1993), and calendar management (Dent *et al.* 1992).

An important form of knowledge that can be acquired by observing users' decisions is knowledge of users' preferences. In configuration tasks such as design or scheduling, for example, there may be numerous configurations that satisfy all applicable hard constraints. Users may nevertheless strongly prefer some configurations to others. For example, in the domain of scheduling ground-based telescope observations, there are typically many different schedules that satisfy all hard constraints (such as not pointing the telescope at the sun or below the horizon, not scheduling two observations at the same time, *etc.*). However, such schedules may differ significantly in factors such as the airmass[1] and research priority of each scheduled observation. Choosing among such schedules requires a model of the relative desirability of schedules as a function of their relevant attributes.

Knowledge of users' preferences can be expressed as a *preference predicate* (Utgoff & Saxena 1987) $P_Q(x, y) \equiv [Q(x) > Q(y)]$, where $Q(s)$ is an evaluation function that expresses the "quality" of state s. A learning apprentice can acquire a user's criteria for the relative desirability of alternative states by learning a preference predicate P_Q from a set of training instances $P_Q(s_i, s_j)$ produced by the user during normal use of the system. For example, each time a learning apprentice suggests a state s_1 and the user rejects s_1 in favor of some other state s_2, the apprentice has an opportunity to acquire the training instance $P_Q(s_2, s_1)$.

Our interest in acquisition of preference predicates arose from a project to develop an intelligent assistant for scheduling ground-based telescope observations, the *Observing Assistant* (OA)(Broos 1993). In developing OA, we found that astronomers could identify the relevant attributes of observation schedules but were typically unable to articulate general criteria for preferring one schedule over another in terms of these attributes. Moreover, it appeared that individual astronomers often differ significantly in their preferences. These factors cast doubt on the feasibility of devising an *a priori* evaluation function appropriate for multiple users. A more promising approach was to develop a learning apprentice system capable of forming "personalized knowledge-based systems" (Dent *et al.* 1992) by acquiring the scheduling preferences of individual astronomers.

The next section describes previous approaches to the problem of acquiring preference criteria and proposes a novel algorithm for this task called *Compositional Instance-Based Learning* (CIBL). Section three describes a series of experiments comparing the performance of CIBL to that of alternate approaches, both in learning to predict astronomer's actual scheduling preferences and in learning artificial preference criteria.

Algorithms for Learning P_Q

Previous approaches to acquisition of preference predicates from sets of training instances have used inductive learning methods to form generalizations from sets of training instances (Utgoff & Saxena 1987; Utgoff & Clouse 1991). One approach has been to

[1]The airmass of an observation is a measure of the amount of atmosphere between the star and the observer. Airmass can be minimized by observing a star at the time midway between its rising time and setting time.

use decision tree induction algorithms, such as ID3 (Quinlan 1986), to induce a general representation for P_Q. An alternative approach, termed the *state preference method*, uses parameter adjustment to learn a set of feature weights $\mathbf{W}$ such that for every training instance, $P_Q(x, y)$, $\mathbf{W}(\mathbf{F}(x) - \mathbf{F}(y)) > 0$, where $\mathbf{F}(n)$ is a vector of numeric attributes representing state n (Utgoff & Clouse 1991).

However, the complexity of astronomers' explanations for preferring one schedule over another led us to hypothesize that the underlying evaluation function Q for astronomical observation schedules, as with preference predicates in many other domains (Callan, Fawcett, & Rissland 1991), is typically not linear, and that the instances of P_Q are therefore not linearly separable. If correct, this hypothesis would imply that learning algorithms that presuppose linear separability, such as the state preference method, are inappropriate for this domain.

Decision tree induction algorithms such as ID3 are suitable for nonlinearly separable data. However, the performance of decision tree induction algorithms has been shown to be sometimes weaker than that of instance-based algorithms when the training set is sparse or the concept being acquired is highly "polymorphic" (Aha 1992). Our hypothesis concerning the complexity of astronomers' preference predicates suggested that these factors would often characterize acquisition of observation scheduling preference predicates. We therefore turned to an instance-based approach to this problem.

Instance-Based Learning of P_Q

Instance-based learning (IBL) is a strategy in which concepts are represented by exemplars rather than by generalizations induced from those exemplars (Stanfill & Waltz 1986). Perhaps the simplest form of instance-based learning is k-nearest-neighbor (k-NN) classification, which classifies a new instance according to the majority classification of its k nearest neighbors in feature space. A straightforward 1-NN strategy for learning preference predicates, which we term *1ARC*, represents training instances as arcs in feature space. For example, on a two dimensional feature space $S = \Re^2$, the set of training instances $\{P_Q(A, B), P_Q(C, D), P_Q(E, F)\}$ is represented as shown in Figure 1 by the training arcs $\overleftarrow{AB}$, $\overleftarrow{CD}$, and $\overleftarrow{EF}$ (where $\overleftarrow{XY} \equiv P_Q(X, Y)$).

Ranking a new pair of objects, X and Y, is equivalent to determining whether $P_Q(X, Y)$ or $P_Q(Y, X)$ is satisfied. The 1ARC algorithm begins by finding the training arc that best matches the hypothesis $P_Q(X, Y) \equiv \overleftarrow{XY}$. The dissimilarity between $\overleftarrow{XY}$ and a training arc is measured by the sum of the Euclidean distances between (1) Y and the tail of the training arc and (2) X and the head of the training arc. The dissimilarity between $\overleftarrow{XY}$ and the training arc that it matches most closely, *i.e.*, for which the dissimilarity is least, is a measure in the confidence in the hypothesis $P_Q(X, Y)$. In Figure 1, for example, the training arc $\overleftarrow{EF}$ best matches $\overleftarrow{XY}$ with a dissimilarity of $dist(Y, F) + dist(X, E)$ represented by the dotted lines.

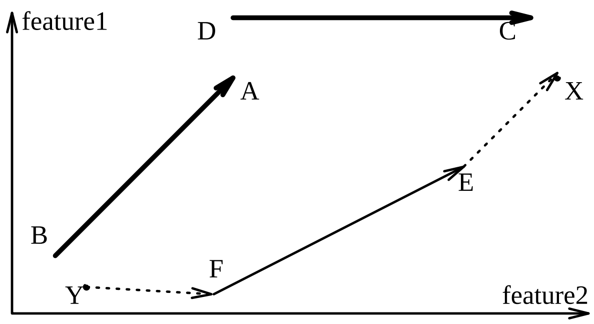

Figure 1: The best match to $\overleftarrow{XY}$ found by 1ARC.

In the same way, 1ARC then finds the best match and confidence measure for the alternate hypothesis $P_Q(Y, X)$. The hypothesis with the strongest measure of confidence determines 1ARC's estimate of the ranking between X and Y. In this case, $\overleftarrow{XY}$ matches training arc $\overleftarrow{EF}$ more strongly than $\overleftarrow{YX}$ matches any training arc, so 1ARC concludes that $P_Q(X, Y)$.

An important limitation of k-NN algorithms, such as 1ARC, is that they are unable to exploit the transitivity of preference predicates. For example, given the situation in Figure 1, it should be possible to conclude $P_Q(X, Y)$ by the reasoning "X is close to C; C is preferred to D; D is close to A; A is preferred to B; B is close to Y". However, the majority vote policy of standard k-NN methods does not permit reasoning involving the serial composition of multiple instances.

Compositional Instance-Based Learning

CIBL (Compositional Instance-Based Learning) is a strategy that permits multiple training instances to be composed to rank a new pair of objects. Like 1ARC, CIBL ranks two new objects, X and Y, by determining whether it has greater confidence in the path from X to Y or in the path from Y to X. CIBL differs from 1ARC in that it can construct a path between two new objects by sequentially connecting multiple training arcs. Such a path seeks to follow a contour of the underlying evaluation function having positive slope.

For example, given the situation shown in Figure 1, CIBL begins by searching for a path from Y to X, supporting the hypothesis $P_Q(X, Y)$. As shown in Figure 2, CIBL forms the *uncertain arcs* U1, U2, and U3. The cost of the path from Y to X is the sum of the Euclidean lengths of the uncertain arcs U1, U2, and U3; the path from the tail to the head of a training arc has zero cost. In a similar fashion, a path is constructed from X to Y. The path with lower cost determines the better estimate of the ranking of X and Y.

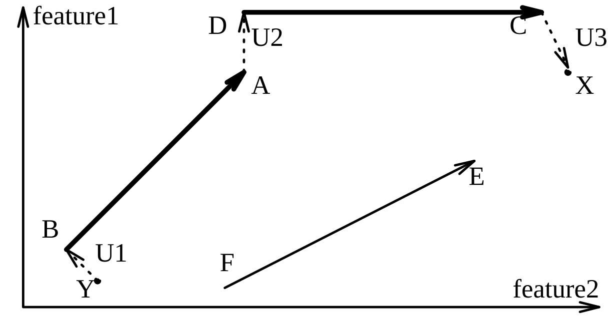

Figure 2: The best match to $\overleftarrow{XY}$ found by CIBL.

In practice, CIBL constructs all possible uncertain arcs, forming a dense graph with two special nodes, X and Y (for clarity, Figure 2 shows only those uncertain arcs on the best path from Y to X). The standard Dijkstra algorithm (Aho, Hopcroft, & Ullman 1974) is then used to find the lowest cost path connecting Y and X, where edges from the tail to the head of a training arc are assigned zero cost and edges representing uncertain arcs are assigned a cost equal to their Euclidean length.

Empirical Evaluation

We conducted three sets of experiments to evaluate the performance of CIBL with respect to the performance of a traditional tree-induction method, ID3, and our 1-NN algorithm, 1ARC. In CIBL, pairs of schedules represented by N attributes were mapped directly onto arcs in an N-dimensional feature space. However, ID3's decision tree model requires that pairs of schedules (arcs) be transformed into positive and negative instances of a concept or class and that concept instances be represented as a single vector of features. We chose to represent arcs using all of their relevant properties: the location of the head, the location of the tail, the normalized direction, and the magnitude. Thus the training set given to ID3 contains, for each training arc $P_Q(X,Y)$, a positive and a negative instance of the concept "Y is preferred to X":

$$< +, Y, X, \frac{(X-Y)}{\|X-Y\|}, \|X-Y\| >$$
$$< -, X, Y, \frac{(Y-X)}{\|X-Y\|}, \|X-Y\| >.$$

Since the (3N+1) elements of each ID3 feature vector are real numbers, we used an implementation of ID3 supplied by Ray Mooney that handles real-valued features in the manner proposed by (Quinlan 1986).

Artificial Domain Experiments

We first compared the accuracy of CIBL to that of 1ARC and ID3 on the task of learning preference functions, P_Q, for a variety of artificial evaluation functions, Q, shown in Figure 3. With the exception of Q_8, all of the evaluation functions were defined on the feature space $S = [0,1] \times [0,1]$. For each Q function, we randomly generated instances of the associated preference predicate, $P_Q(X,Y)$, representing the knowledge "X is preferred over Y" for $X, Y \in S$. Each model was trained on a set of instances of size $\|TS\| \in \{2, 8, 32, 128\}$ and was then tested on a different set of instances of size 1000. Each $< model, Q, \|TS\| >$ triplet was trained and tested four times and an error rate was calculated by counting the incorrect rankings in the four tests.

$Q_1(f_1,f_2)$	$= f_1 + 10f_2$	1x10 plane
$Q_2(f_1,f_2)$	$= (f_1 - \frac{1}{2})^2 + (f_2 - \frac{1}{2})^2$	quadratic
$Q_3(f_1,f_2)$	$= sin(2\pi(f_1 + f_2))$	sinusoid
$Q_4(f_1,f_2)$	$= \begin{cases} f_2 & \text{if } f_1 \leq \frac{1}{2} \\ 1 - f_2 & \text{else} \end{cases}$	crossed planes
$Q_5(f_1,f_2)$	$= f_1 + f_2$	1x1 plane
$Q_6(f_1,f_2)$	$= exp(f_1^2 + f_2^2)$	exponential
$Q_7(f_1,f_2)$	$= \begin{cases} f_1 + f_2 & \text{if } f_1 \leq \frac{1}{2} \\ 1 + f_2 - f_1 & \text{else} \end{cases}$	folded plane
$Q_8(f_1,f_2,f_3)$	$= (f_1 - \frac{1}{2})^2 + (f_2 - \frac{1}{2})^2.$	Q_2 in 3-D

Figure 3: Quality Functions.

Figure 4: Error Rates for Various Q Functions.

For all evaluation functions tested, CIBL had a lower error rate than 1ARC, as shown in Figure 4. This indicates that CIBL's strategy of composing multiple exemplars is superior to 1ARC's traditional 1-NN approach. ID3 performed better than CIBL on the evaluation functions Q_1 (plane), Q_5 (plane), and Q_6 (exponential), which have no change in the sign of their derivative. Both 1ARC and CIBL performed significantly better than ID3 on the evaluation functions Q_2 (quadratic), Q_3 (sinusoidal), and Q_4 (crossed planes), which exhibit changes in the sign of their derivatives in the form of local extrema or a discontinuity. These data indicate that CIBL generally performs better than ID3 when the evaluation function is "complex" in the sense of containing a local extremum or discontinuity. However, ID3 and CIBL performed equally on Q_7

(folded plane), a function with an abrupt change in the sign of its derivative.

As expected, the addition of an irrelevant feature to the feature space—a feature that has no effect on the evaluation function—did not affect ID3's performance. However, because CIBL's Euclidean distance metric, used to assign costs to uncertain arcs, counts all features equally, CIBL's accuracy was degraded by the addition of an irrelevant feature, as shown in the testing data for Q_8 (the same 2-D quadratic as function Q_2 with an irrelevant third feature added). The sensitivity to irrelevant features exhibited by CIBL has been observed in other studies of instance-based learning (Aha 1989).

The second experiment tested the hypothesis that the relative performance of CIBL and ID3 depends on the dimensionality of the feature space as well as on the complexity of the quality function underlying the preference predicate. To test this hypothesis, we compared the ability of CIBL and ID3 to acquire a preference predicate for a linear Q on artificial data of dimensionality 2, 3, 5 and 10. As with the first experiment, both training and testing instances were uniformly distributed through the feature space. The results, set forth in Figure 5, show that for linear Q ID3 has a lower error rate in feature spaces of dimensionality less than 5, the error rate is comparable for dimensionality equal to 5, and CIBL has a lower error rate in feature spaces of dimensionality greater than 5.

Figure 5: Cumulative error rate of CIBL and ID3 for linear Q in feature spaces of dimension 2, 3, 5, and 10.

Scheduling Experiments with Astronomers

Observing Assistant The Observing Assistant is a decision support system to assist astronomers in scheduling ground-based telescope observations. OA acts as a smart schedule editor that assists in the incremental process of schedule construction used by astronomers to construct schedules by hand: starting with an empty schedule, OA suggests refinements to the current partial schedule by adding one object from the astronomer's catalog of desired observations.

OA uses its model of the astronomer's preference predicate to sort the refinements of a partial schedule S. The set of refinements of S consists of each placement into S of an unscheduled object from the astronomer's catalog that results in a new schedule satisfying all hard constraints. The highest ranked refinement is then suggested to the user. If the user rejects the proposed refinement s_i in favor of some other refinement s_j, OA acquires the training instance $P_Q(s_j, s_i)$. A separate model is maintained for each astronomer.[2]

Interactive Learning Experiment The artificial domain experiments indicated that CIBL always performs at least as well as 1ARC and that the relative performance of CIBL and ID3 depends upon the nature of the underlying quality function Q and the dimensionality of the feature space. The second set of experiments compared the relative effectiveness of CIBL to that of ID3 on the task of learning an astronomer's scheduling behavior in the context of the Observing Assistant. Two different versions of OA were implemented: OA-CIBL used the CIBL learning method; and OA-ID3 used the ID3 learning method.

A typical observing catalog of astronomical objects was provided by the director of a ground-based observatory. An astronomer at the observatory scheduled this catalog twice, once using OA-CIBL and once using OA-ID3. The catalog comprised three nights of observations, so a total of six nights were scheduled (three nights per catalog, two different learning methods). The six learning sessions were interleaved so that the astronomer did not know which learning method was in use. Each time the astronomer made a ranking decision, that is, each time the astronomer expressed a preference for a particular schedule in a set of schedules, data were collected on OA's performance.

The relative performance of the learning algorithms was measured in two different ways. The first measure was cumulative error rate, which indicates how often each model failed to identify correctly the astronomer's preferred schedule. The cumulative error rate of CIBL was significantly lower (37%) than that of ID3 (47%), indicating that the astronomer accepted CIBL's suggested refinement more often than she accepted ID3's.

The second measure of performance was a *linear payout metric* under which a model is rewarded by $\frac{2(n-m)}{n-1} - 1$ if the model assigned the user's first choice

[2]The preference model of the initial implementation of OA uses the following attributes of observing schedules:

- the priority of the observation most recently added to the schedule
- the duration of the most recent observation
- the maximum airmass of the most recent observation
- the optimal airmass of the most recent observation, *i.e.,* the lowest airmass it achieves during the entire night
- the average airmass of the other objects in the schedule.

Several additional attributes, such as total telescope slew time, would need to be added for a complete model of the factors considered by astronomers in scheduling.

out of n objects a rank of m. This metric rewards a scheduler by $+1.0$ when the user's chosen schedule was ranked first and by -1.0 when the user's chosen schedule was ranked last. The expected value of a preference predicate model with no knowledge is zero. Figure 6 shows the cumulative payout data for OA-CIBL and OA-ID3, indicating that both had about the same ability to predict the astronomer's behavior. The relatively high payout from both methods—over 40 after 62 instances—indicates that both methods rapidly acquired a sufficiently accurate preference model to provide useful advice to the astronomer.[3]

Figure 6: Cumulative payout. The 45° line represents the cumulative payout of a perfect model of the astronomer's preference predicate.

Replay Experiments In addition to directly measuring the relative performance of CIBL and ID3 as the learning component of OA, the learning methods were compared on two sets of approximately 135 preference instances recorded from each of two different astronomers who used OA to schedule 6 nights of observations.

In the first experiment, each astronomer's preference instances were used to train CIBL and ID3 separately using a learn-on-failure protocol. The two states contained in each preference instance were given to the model (CIBL or ID3) for ranking, and the model learned the instance only if it ranked the states incorrectly. Both models had approximately equal cumulative error rates (astronomer #1: CIBL-21%, ID3-22%; astronomer #2: CIBL-14%, ID3-21%), confirming the result of the interactive learning experiment described above that CIBL and ID3 had comparable abilities to predict astronomer's behavior.

The second replay experiment tested the hypothesis that different astronomers use distinct preference predicates. The two sets of preference instances were each randomly partitioned into two subsets. One partition was used to train a preference predicate model. The model's error rate was then measured on the task of predicting the preferences contained in the other partitions. This experiment was performed under twelve different configurations to cover all the possible permutations of three configuration variables: the preference model used (CIBL or ID3); the source of the training partition (astronomer #1 or astronomer #2); and the size of the training partition (45, 68, or 90 instances). The experiment was repeated 10 times for each testing configuration.

Over all 120 tests, the average error rate for ranking instances from the set used to train the model (8.7%) was significantly lower than the average error rate for ranking instances from the other astronomer's set (25.0%). This indicates that there was a significant difference between the scheduling behaviors of the two astronomers we tested, confirming the hypothesis that different astronomers require different preference models.

Scheduling Experiments With Artificial P_Q

The final experiment tested whether the dependence of the relative performance of CIBL and ID3 on the complexity of the underlying quality function Q, which was observed in an artificial domain, also applies when scheduling actual astronomical observations. To test this hypothesis, OA-CIBL and OA-ID3 were rerun on the catalog of observations using each of the quality functions set forth in Figure 7 as an oracle in place of a human astronomer. As shown in Figure 8, the results confirmed that CIBL's performance relative to ID3 improves with increasingly complex Q: ID3 is more accurate than CIBL for linear Q,[4] CIBL is slightly more accurate for quadratic Q, and CIBL is much more accurate for sinusoid Q.

$$
\begin{array}{ll}
Q_9 = f_1 - f_2 + f_3 + f_4 - f_5 & \text{plane} \\
Q_{10} = -[(f_1-1)^2 + (f_2-2)^2 + (f_3-\frac{3}{2})^2 & \\
\quad + (f_4-\frac{1}{2})^2 + (f_5-2)^2] & \text{quadratic} \\
Q_{11} = sin(\pi\sqrt{f_1^2+f_2^2+f_3^2+f_4^2+f_5^2}) & \text{sinusoid}
\end{array}
$$

Figure 7: 5-D Quality Functions.

Conclusion

Learning apprentice acquisition of preference predicates, as typified by OA-CIBL and OA-ID3, is appropriate when (1) users can identify the relevant characteristics of problem-solving states, (2) these state

[3] The slightly higher payout for OA-ID3, notwithstanding its somewhat lower accuracy, indicates that the average magnitude of errors was somewhat greater for OA-CIBL.

[4] This result appears to be inconsistent with the second artificial domain experiment, in which ID3 and CIBL had comparable accuracy for linear Q in a five-dimensional feature space. However, this disparity is attributable to the differences between the two experiments: (1) the instances used for training and testing were random points in feature space for the earlier experiment but were actual schedules for the later experiment and (2) the task in the earlier experiment was to establish a binary ranking whereas the task in the later experiment was to order a full set of schedule refinements.

Figure 8: Cumulative payout of CIBL and ID3 with 5-D plane, quadratic, sinusoid functions replacing the human astronomer.

characteristics can be adequately represented as an attribute vector, but (3) users differ as to or are unable to articulate evaluation criteria for problem solving states in terms of these attributes.

The empirical evaluation showed that CIBL's strategy of composing instances of preference predicates is superior to a 1-NN instance-based learning strategy unable to compose instances. The relative performance of CIBL and decision tree induction for preference predicate acquisition depended upon (1) the complexity of the preference predicate P_Q being acquired as measured by the underlying evaluation function Q and (2) the dimensionality of the feature space. Irrelevant attributes degraded the performance of CIBL but not ID3. CIBL and ID3 performed comparably when tested as the learning component of a learning apprentice used by an astronomer for scheduling astronomical observations having five real-valued attributes. A replay experiment confirmed the hypothesis that astronomers may differ widely in their scheduling preference predicates.

The empirical evaluation suggests that CIBL is preferable to ID3 as the learning component of a learning apprentice system if representation of the relevant characteristics of problem-solving states requires more than five attributes or if attributes interact in complex ways (*i.e.*, the underlying quality function has extrema or discontinuities), provided that all attributes are relevant. Conversely, ID3 is preferable if there are fewer than five attributes and the attributes do not interact in a complex fashion (*i.e.*, the quality function has no extrema or discontinuities) or if there are irrelevant attributes.

Acknowledgements

This research was supported in part by the Wyoming Planetary and Space Science Center. We would like to thank Earl Spillar and Leisa Townsley for their assistance with telescope scheduling, Peter Turney for his suggestions for applying ID3 to real-valued data, and Ray Mooney for his implementation of ID3.

References

Aha, D. 1989. Incremental, instance-based learning of independent and graded concepts. In *Proceedings of the Sixth International Workshop on Machine Learning*, 387–391.

Aha, D. 1992. Generalizing from case studies: A case study. In *Proceedings of the Ninth International Workshop on Machine Learning*, 1–10.

Aho, A.; Hopcroft, J.; and Ullman, J. 1974. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Publishing Co.

Broos, P. 1993. An expert system for telescope scheduling. Master's thesis, University of Wyoming.

Callan, J.; Fawcett, T.; and Rissland, E. 1991. Adaptive case-based reasoning. In *Proceedings of the Third DARPA Case-Based Reasoning Workshop*, 179–190. Morgan Kaufmann.

Dent, L.; Boticario, J.; McDermott, J.; Mitchell, T.; and Zabowski, D. 1992. A personal learning apprentice. In *Proceedings of Tenth National Conference on Artificial Intelligence*, 96–103. San Jose, CA: AAAI Press/MIT Press.

Maes, P., and Kozierok, R. 1993. Learning interface agents. In *Proceedings of Eleventh National Conference on Artificial Intelligence*, 459–466. Washington, D.C.: AAAI Press/MIT Press.

Mahadevan, S.; Mitchell, T.; Mostow, J.; Steinberg, L.; and Tadepalli, P. 1993. An apprentice-based approach to knowledge acquisition. *Artificial Intelligence* 64(1).

Mitchell, T.; Mahadevan, S.; and Steinberg, L. 1985. Leap: A learning apprentice for vlsi design. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann.

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

Stanfill, C., and Waltz, D. 1986. Toward memory-based reasoning. *Communications of the ACM* 29(12).

Utgoff, P., and Clouse. 1991. Two kinds of training information for evaluation function learning. In *Proceedings of Ninth National Conference on Artificial Intelligence*, 596–600. Anaheim: AAAI Press/MIT Press.

Utgoff, P., and Saxena, S. 1987. Learning a preference predicate. In *Proceedings of the Fourth International Workshop on Machine Learning*, 115–121.