

Experience with Rule Induction and k-Nearest Neighbor Methods for Interface Agents that Learn

Terry R. Payne, Peter Edwards, and Claire L. Green

Abstract—Interface agents are being developed to assist users with a variety of tasks. To perform effectively, such agents need knowledge of user preferences. An agent architecture has been developed which observes a user performing tasks, and identifies features which can be used as training data by a learning algorithm. Using the learned profile, an agent can give advice to the user on dealing with new situations. The architecture has been applied to two different information filtering domains: classifying incoming mail messages (Magi) and identifying interesting USENet news articles (UNA). This paper describes the architecture and examines the results of experimentation with different learning algorithms and different feature extraction strategies within these domains.

Index Terms—Machine learning, interface agent, information filtering, intelligent USENet news reader, intelligent e-mail filter, agent architecture, instance-based learning, rule induction.

1 INTRODUCTION

IN recent years, interface agents have been developed to assist users with tasks such as filtering incoming information [1], [2] and arranging diary appointments [3], [4]. If the agent is to be of assistance, it requires knowledge about the domain application and/or the user. Many agents employ learning techniques to acquire this knowledge. They may learn autonomously by observing the user or be explicitly trained.

Traditionally, two approaches have been used to provide an agent with knowledge about a task domain. The first and most common approach is for users to provide their own rules. This is used in systems such as *Oval* [5] which determines if a mail message is of interest, and if so, what actions should be performed. Systems such as this utilize a scripting language to specify rules. However, learning and applying this scripting language may discourage nontechnical users from using the system. As well as understanding exactly how they require the system to behave, a user must appreciate how the agent will perform with the rules. The user also has to be responsible for maintaining the rules over time as their interests change.

The second approach makes use of traditional knowledge engineering techniques to identify background knowledge about the application and the user (such as the Unix consultant UCEgo [6]). While this shifts the task of programming the agent from the user to the Knowledge Engineer, the agent will not be customized for a particular user. Thus, the approach cannot be used for personalized tasks such as information filtering.

The use of learning techniques to develop a *profile* of an individual user's preferences not only eliminates the need for programming rules, but also allows the agent to adapt to changes. There are many ways in which a system can learn from the user [7]. Approaches such as programming by demonstration provide good training examples. In comparison, apprentice sys-

tems acquire knowledge by observing and analyzing the user's behavior.

The profile should be used by the agent to decide what assistance to provide and to determine some measure of its reliability. This can be achieved by generating a *confidence rating*, which provides an indication of the agent's confidence in its predictions. Unless the advice generated is accurate, the user will fail to trust the agent. Thus, the user should be able to override agent decisions if necessary. This feedback can be used to refine the profile, identifying aspects which result in poor behavior and promoting those which improve it.

The work described here details an interface agent architecture which learns from observations, and describes how it has been applied to two different information filtering domains; that of classifying incoming mail messages (Magi) and identifying interesting USENet news articles (UNA).

2 LEARNING AND INFORMATION FILTERING

A number of applications have been developed which employ machine learning techniques to assist a user with filtering USENet news articles and e-mail messages. Lang [8] has developed a news-filtering system (NewsWeeder) which learns a user profile by allowing the user to rate their interest in each article on a scale of 1-5. The system currently focuses on using the content of the article to determine its relevance, which is known as *content-based filtering*. As a user reads each article, they provide a rating. Each night, the system generates a new profile based on these ratings. As well as reading different newsgroups, a user can also read NewsWeeder's *virtual newsgroups*. Such a newsgroup contains a personalized, sorted list of the top rated articles determined by the agent.

Lang has explored two different strategies to identify articles of interest. A popular technique used in Information Retrieval called *term-frequency/inverse-document frequency* weighting (TF-IDF) [9] has been used as a benchmark from which to compare another technique, the *Minimum Description Length* principal (MDL), described in [8]. Both strategies rely on breaking up the article body into *tokens* and counting their occurrence to create a vector of token counts.

NewT (News Tailor) [1] adopts a genetic algorithm based approach to identify articles of interest to the user. A set of profiles are applied to new articles to identify those the user would find interesting. Each profile has a rating, which measures how effective it is in identifying such articles. Genetic processes, such as *crossover* and *mutation* are used to create new profiles which may outperform current profiles, or explore new areas of interest. Profiles with a low rating are eliminated, ensuring that the number of profiles is kept to a manageable level.

NewT periodically filters new articles to determine which ones could be of interest to the user. This is done by converting articles into their vector space representation [9]. Each article is tested against the profiles, and ranked according to the closeness of the match. The highest ranking articles are then presented to the user. The user provides positive or negative feedback as the articles are browsed, which is then reflected by changes in the rating of a profile.

A different approach has been used in the development of the mail filtering agent, Maxims [10]. An earlier calendar management system [4] was adapted to produce a generic agent architecture which could be attached to any application. Maxims learns to prioritize, delete, forward, sort, and archive mail messages on behalf of the user. The agent uses the sender and recipient fields of a message (including `cc`: recipients), and keywords from the subject field. Other information such as whether the message has been read, whether it is a reply to a previous message, etc., is also used.

• The authors are with the Department of Computing Science, King's College, University of Aberdeen, Aberdeen, Scotland, AB24 3UE. E-mail {terry, pedwards, clg}@csd.abdn.ac.uk.

Manuscript received Sept. 12, 1996.

For information on obtaining reprints of this article, please send e-mail to: transkde@computer.org, and reference IEEECS Log Number K97019.

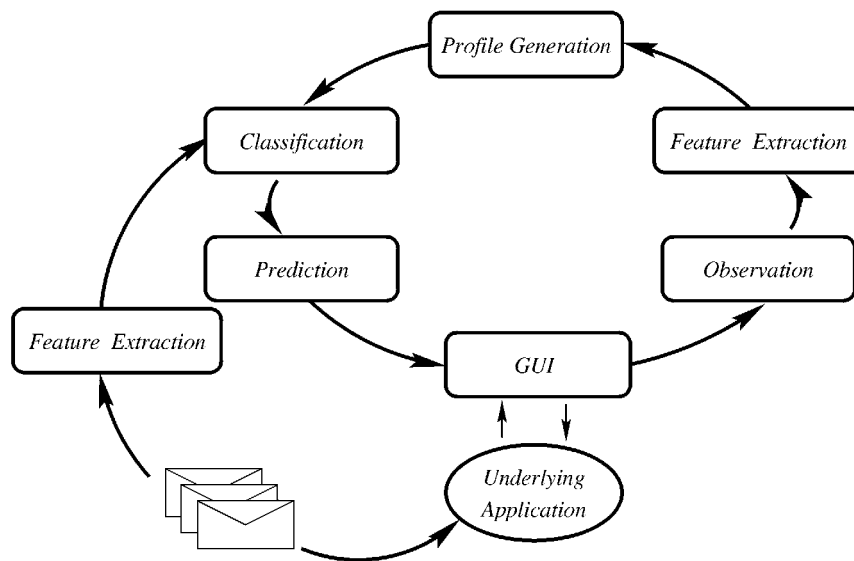


Fig. 1. A learning interface agent architecture.

As the user reads their mail and performs certain actions, the situation is recorded. When new messages arrive, the agent uses *Memory-Based Reasoning* [11] to find the closest stored situation, and make a recommendation to the user. Confidence ratings are generated and compared to two threshold values: a *tell-me* threshold and a *do-it* threshold. These are used by the agent to determine whether to simply recommend actions or whether to act on behalf of the user.

A caricature, in the form of a face, is used to indicate the agent's confidence in its recommendations. Different expressions are used, so the user knows if an action has been performed on their behalf, or if the agent wants to advise them. Depending on whether the user confirms or rejects the recommendation, the agent will question the user to determine what factors were important in this response. This feedback is used to adjust priority weights which affect future recommendations.

Other agents have been developed which employ machine learning techniques to assist users exploring the World Wide Web, such as WebWatcher [12].

3 AGENT MODEL

Our agent architecture¹ is shown in Fig. 1. A graphical user interface (GUI) is used to interact with the underlying application. As it is used, observations are made from which the agent can induce a user profile. These observations, consisting of articles and actions performed on them, are used to generate training examples, by passing them to the feature extraction module. The training examples are then used by a learning algorithm to induce a user profile. New articles are also processed by the feature extraction module and output passed to the classification stage. The user profile is used to generate a classification such as a mail processing action (Magi), or an interest rating (UNA). The resulting classifications are evaluated by the prediction stage, and a prediction is made.

The feature extraction module identifies fields in the articles such as the *author* or the *subject*, and extracts values from them. Words are also extracted from the article body based on how frequently they occur within the text. Other information such as the length of the article can also be determined. Within an article, each field can gen-

erate a number of values. These values are used to generate the user profile and subsequently to make predictions on new articles.

Two different learning paradigms have been explored within this architecture: a rule induction algorithm, CN2 [14], and a *k*-nearest neighbor algorithm (*k*-NN) IBPL [13]. The initial motivation for using CN2 was that this algorithm generated human comprehensible rules by performing induction over training examples containing specific features [2]. IBPL was explored to contrast with the symbolic approach, and to overcome some of the problems encountered by CN2 when learning from textual data.

Fig. 2 shows how a user profile is generated by both learning algorithms from a single mail message. The feature extraction module generates feature sets for each of the fields within the mail message. These feature sets are then mapped to training examples. The exact format of the examples used depends on the nature of the learning algorithm. Many examples are generated for CN2 as it expects single values for each attribute. These are then used to induce ordered 'if-then' (production) rules. A single example is generated for IBPL, as it has been designed to learn directly from the feature sets.

CN2 is a supervised learning algorithm that constructs ordered production rules from a set of preclassified examples. It performs a 'best-set-so-far' beam search on a size limited set of *complexes*, where each complex is a conjunction of attribute tests associated with a class. Each complex considered is specialized to maximize the number of examples it covers from its class, while reducing the number of examples covered from other classes. These complexes are then combined to produce the resulting set of rules. See [14] for a complete discussion of the algorithm.

When using the production rules for the classification stage, many examples are generated from each new article. This provides a means of producing a confidence rating for each prediction made by the agent. The examples may fire different rules, leading to different classifications. The number of rules which fire for each classification are therefore summed, and a confidence rating is generated.

As mentioned above, the IBPL *k*-NN algorithm was developed to learn from sets of values. Nearest-neighbor algorithms derive a classification by comparing a new instance with previously classified instances [15]. It is possible to modify the comparison so that multiple values can be compared for each attribute. The value-distance metric, used in the memory-based reasoning algorithm [11], provides a means of comparing two symbolic instances by

1. A complete description of the architecture and the learning algorithms described above can be found in [13]. Also included is a detailed evaluation of both algorithms when applied to classifying mail messages as part of the Magi system.

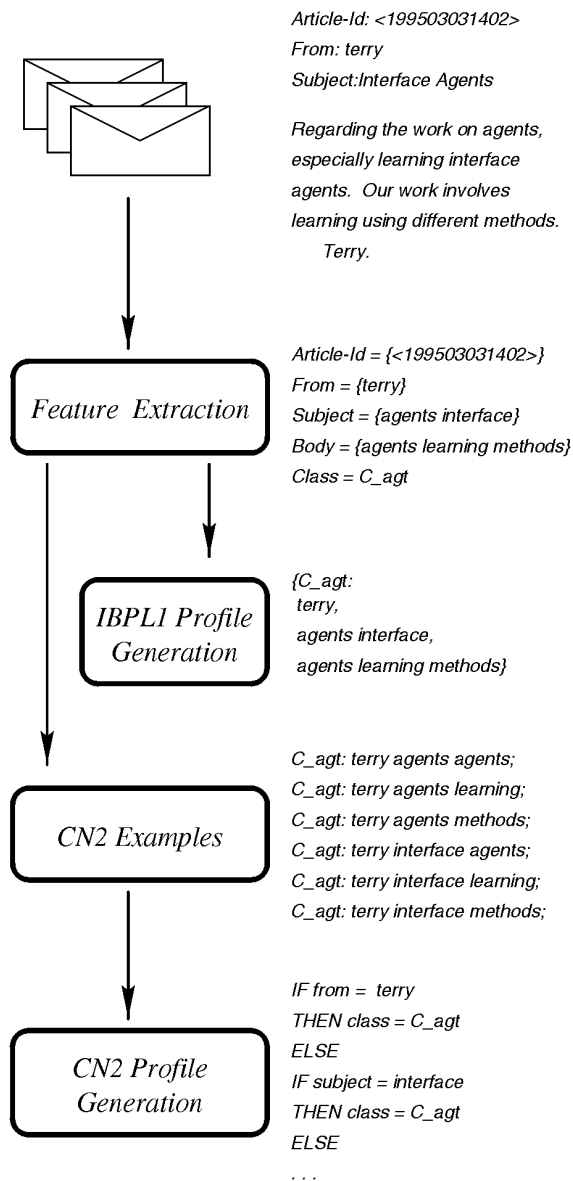


Fig. 2. Generating a profile using CN2 or IBPL.

calculating a similarity measure (*distance*) between values. This is done by mapping each symbol to a distribution matrix. For this reason, this metric was modified (see Equation 1) so that multiple distances could be calculated and averaged when comparing attributes. This resulted in the algorithm IBPL [13].

$$\Delta(\tau, \rho) = \sum_{f \in \mathcal{F}_\tau} \left(\frac{\sum_{i \in \tau, f} \sum_{j \in \rho, f} d(i, j) w(i)}{|\tau, f| \times |\rho, f|} \right) \quad (1)$$

$$d(i, j) = \sum_{c \in C} (\delta(c, i) - \delta(c, j))^2 \quad (2)$$

$$w(i) = \sqrt{\sum_{c \in C} \delta(c, i)^2} \quad (3)$$

Each article generates a single training episode, in which each attribute contains a set of all the extracted values from the relevant field. A similarity measure is calculated by determining the distance $d(i, j)$ between two values for feature set f (2), and a

weighting value $w(i)$ for value i (3). The k most similar instances are found and used to determine the final classification. A confidence rating is also generated as part of this classification process. Table 1 summarizes the notation used.

TABLE 1
 NOTATION USED IN EQUATIONS (1), (2), AND (3)

Symbol	Description
C	The set of all classes in the training set.
\mathcal{F}	Set of attributes.
τ	Unclassified instance.
ρ	Instance in training set.
f	A single feature set.
i	A value from a feature set in τ .
j	A value from a feature set in ρ .
$ \rho, f $	The number of values within the feature set f in instance ρ .
$\delta(c, x)$	Ratio of the number of times value x occurs in training instances of class c , to the number of times x occurs in the training set.

As new observations are made and new training examples created, the examples are time stamped. This way the training set can be pruned with respect to time. In addition to reducing the number of training examples used, this also removes old examples which may not reflect changes in user behavior.

The two learning approaches differ in the way new training examples are integrated into the user profile. CN2 periodically (e.g., every night) induces a new rule set, pruning out old examples and adding new examples. Hence new examples do not affect performance until new rules are induced. IBPL, however, introduces new examples into the training set as soon as they are created and, thus, the effects are immediate.

We will now describe two agents, Magi and UNA, which are based on the agent architecture shown in Fig. 1. The use of both learning algorithms has been explored with these systems and is described later.

4 MAGI

Magi aids a user in sorting incoming electronic mail [2], [13]. In essence, the system is an apprentice which autonomously observes and analyzes user behavior in dealing with mail. By interacting with a modified version of *Xmail*, a user can send and read mail messages, and organize their mail box. For each session, a *session logfile* is created which contains the user's actions and the messages on which the actions were performed.

Periodically, features are extracted from the messages in this logfile; it is these which are used to generate the user profile. Features extracted from incoming mail messages are tested by the classification engine, and a confidence rating is generated. The prediction is considered valid if its confidence rating is greater than a lower threshold value, known as the *predictive threshold*. Valid predictions are stored by the agent for presentation to the user.

The user is informed if new messages have arrived when they next use the application. At this point the user can instruct the agent to perform its suggested actions on the messages, or can browse the predictions. For each type of action that can be predicted, there exists a *confidence threshold*. Only actions with ratings greater than this will be invoked. The rationale behind this is that certain predicted actions, if incorrect, can be tolerated, such as storing a message in the wrong mailbox. However, actions such as those that delete mail or forward messages to other recipients are more critical. Thus, the agent requires a higher level of certainty in such predictions before performing them. This second threshold value is therefore used to determine those predictions which require user confirmation before being performed.

A prediction browser allows the user to monitor actions sug-

gested by the agent and, thus, establish trust in the agent's predictive ability. The browser displays a summary of the predicted actions and indicates those with a sufficiently high confidence rating to be invoked. The user can either confirm predictions with low ratings, or reject highly rated predictions, thus overriding the agent's decision. This feedback could be used to adjust the confidence threshold for each class of action, such as deletion, etc.²

5 UNA

UNA aids a user by identifying interesting USENet news articles [16]. As the user reads each news article in a modified version of the *xrn* browser, they provide a rating in order to indicate their level of interest in the article. The interest rating is an integer in the range 1-6, conveyed by pressing one of six buttons on the user interface (see Fig. 3). A rating of 1 indicates that the user found the article extremely dull or uninteresting, while a rating of six indicates to the agent that the user found the article highly interesting. Article details and ratings are appended to a *session logfile*. When the user exits the user interface, features are extracted from these observations, and are used to generate the user profile. The user profile is utilized by the classification engine in order to classify future news articles.

Periodically (e.g., every hour) a daemon runs, which identifies the newsgroups the user is subscribed to, and queries the news server to retrieve all new articles posted to each subscribed newsgroup. Features are extracted from these new articles in the same way as for the training data. The articles are then classified and the results passed to the prediction stage, which interprets the results of classification, generating a prediction (on the scale 1-6) of the user's interest in the article.

When the user next reads news, they can choose one of two modes: agent mode or browse mode. When in browse mode, there is no agent intervention in the presentation of articles to the user; all articles are presented, regardless of whether the agent has judged them to be of interest or not. When in agent mode, however, the agent marks any new articles that it has predicted as uninteresting (i.e., given a rating 1-3) as having already been read. All articles which have been predicted as interesting (i.e., given a rating 4-6) or those for which the agent was unable to generate a prediction, are left as unread. With this method, articles believed to be of little or no interest are filtered out.

An agent status window runs permanently in the background of the user's desktop (see Fig. 3). This is a graphical representation of the status of the agent. The user, by simply glancing at the agent status window, can see if any new articles have been received, and if so, whether any of these articles have been classified as interesting. The status window can represent four agent states: *idle*, *learning*, *dull*, or *excited*. The agent is deemed to be *idle* if the daemon is not running, and no new articles have been posted since the user last read news. The agent is *learning* if a profile is being generated from user observations. The *dull* icon indicates that new articles have been posted to at least one newsgroup and that the articles have been classified as uninteresting, whereas when the *excited* icon appears, some interesting articles have been detected.

6 EXPERIMENTATION AND EVALUATION

Experimentation has been performed to compare the performance of both learning algorithms in making accurate predictions for new messages/articles. The Magi test set consisted of 408 mail messages, sorted into 12 classifications, whereas the UNA set contained 1,200 articles split evenly across six newsgroups.

2. The use of feedback to adjust the different confidence thresholds for each type of action has yet to be implemented.

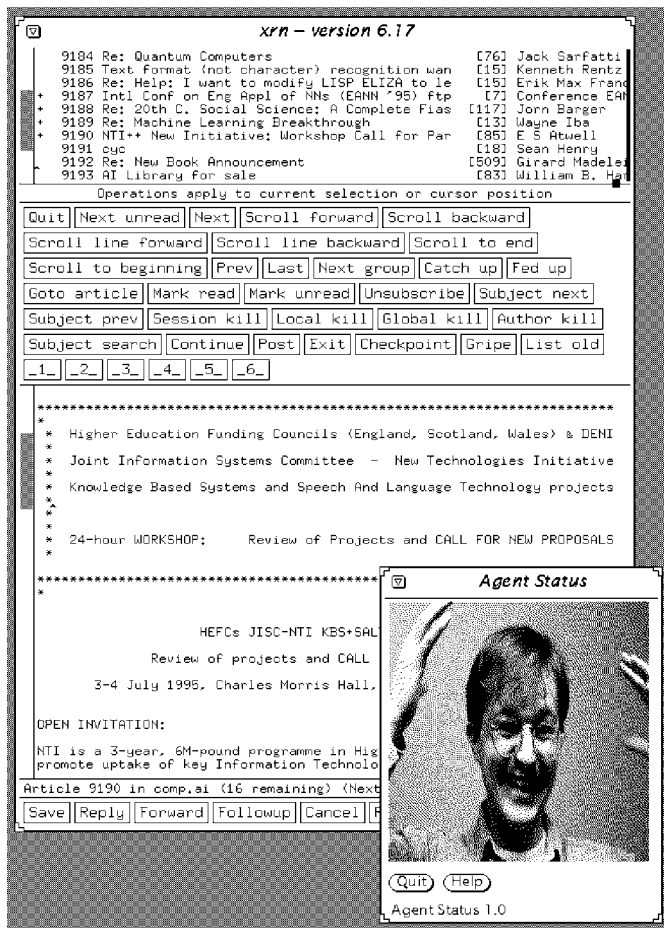


Fig. 3. The UNA user interface.

UNA was tested by rating messages as either interesting or dull, and by providing an integer rating between 1 (dull) and 6 (interesting).

Each Magi classification determined in which mailbox a message should be placed. Mailboxes such as *agents* and *cure* contained messages from a mailing list. The three mailboxes *dai*, *kdd*, and *mead* all contained messages from mailing list digests. The messages in these mailboxes contain individual messages which have been grouped together by a moderator (either manually or automatically). Thus, the message body features selected represent the digest, as opposed to any individual topic within the digest. This data is described fully in [13].

The evaluation was performed on a Sun SPARCserver 1000. Both learning algorithms were implemented in C and were used within a testbed implemented in both C and Bourne shell script.

6.1 Magi Experimentation

The average accuracy of predictions over all mailboxes was higher for CN2 (65 percent) than for IBPL (57 percent). Figs. 4 and 5 show examples of the performance for individual mailboxes: the *agents* mailbox and a small digest mailing list, *dai*. For the digest mailing lists which contained large numbers of messages (27 messages in *kdd*, 90 messages in *mead*) both CN2 and IBPL were able to predict user actions with near 100-percent accuracy. While CN2 produced accurate predictions for the small digest mailbox *dai*, IBPL performed badly (see Fig. 5). This can be explained by considering the voting strategy used by IBPL, where the top k messages are considered in determining the classification. In this case, k was set to 10. Cover [17] demonstrated that a larger value of k results in an

improvement in the behavior of large samples, at the expense of small sample behavior.

Fig. 6 shows the proportion of rules generated for each mailbox by CN2. The digest mailboxes rely on a very small number of rules. This is due to the *From* and *Subject* fields being similar across all messages within a digest. However, for other mailboxes which have different *From* and *Subject* features, there is a steady increase in the number of rules as the number of messages in the mailbox increases. This indicates that there are few features which are common to all messages within a given mailbox. As a result, a larger number of rules are need to cover the larger number of less common features found in the messages.

6.2 UNA Experimentation

Experimentation with UNA concentrated on examining the feature extraction mechanism and investigating whether a hotlist of words would improve performance. The percentage of correctly predicted interest ratings varied between 30 percent and 80 percent, depending on the newsgroup and learning algorithm. With a broad classification (where articles are noted as either interesting or dull), an average of 59 percent of the predictions were correct with CN2, compared to 51 percent when using IBPL. This contrasts with using narrow classifications (six classes—three positive and three negative) where only an average of 27 percent of article ratings were correct with CN2, and 25 percent with IBPL. See Figs. 7 and 8 for examples of UNA predictions for broad and narrow classifications.

Experimentation was also carried out to investigate alternative methods of feature extraction. These methods consisted of extracting the contents of the *From* field, determining the length of the article, and also implementing a user defined hotlist of significant words to identify relevant features, as well as investigating combinations of these methods.

The results were inconclusive for all alternative methods of feature extraction, as none of the methods resulted in a significant improvement in accuracy. For example, for some newsgroups the addition of a hotlist improved performance, while degrading it for others. Figs. 9 and 10 show the change in performance when the basic feature extraction mechanism was adapted to include a hotlist for the newsgroup *sci.psychology*.

It can be seen that there is no significant increase in accuracy with the addition of a hotlist, regardless of the learning algorithm used. As the articles rated by the user are already clustered into a particular newsgroup, the frequently occurring words in one article are likely to occur in many articles across the newsgroup.

7 DISCUSSION

Previous work in this area has concentrated on issues such as the interaction between the user and the agent (e.g., Maxims, NewT), or different classification mechanisms, as in NewsWeeder. The motivation behind our agent architecture was the development of a testbed to explore different aspects of interface agent technology. For example, different learning techniques can be compared and various approaches to feature extraction can be explored.

The two agent systems described in this paper use the same feature extraction mechanism, which extracts words according to word frequency. The underlying assumption here is that words which act as good classifiers for identifying message topics appear frequently. While this model appears to work for Magi, where the task is primarily that of grouping together related messages, it is unsuitable for UNA where articles have already been sorted into topics, or newsgroups.

The performance of UNA degrades significantly when multiple narrow classifications are used. As the number of classes increases, there is a greater chance of features appearing in more than one

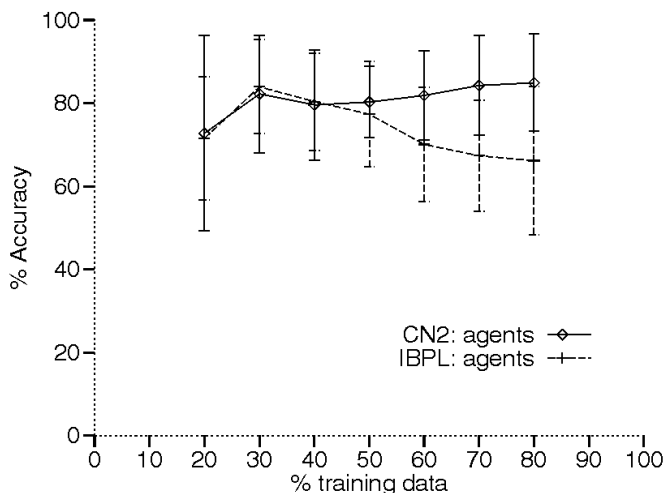


Fig. 4. Accuracy of predictions made for the agents mailbox.




Fig. 5. Accuracy of predictions made for the dai mailbox.

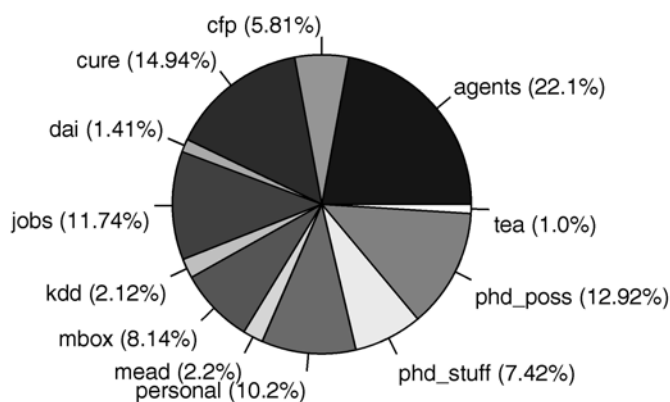


Fig. 6. Percentage of rules covered by Magi for the different mailboxes.

class. Algorithms such as CN2 and IBPL consider each classification as distinct from the others, as a result, such features will be considered as poor classifiers.

An important difference between the two algorithms is the time taken to induce and apply user profiles to new articles. The in-

stance based approach builds a sub-symbolic representation in the form of weights and distances. Unlike rule induction in CN2, these calculations do not involve searching through a large space of possible solutions. The search performed by CN2 is compounded by the large number of features generated by the article body. It was found that tests involving CN2 took significantly (30 to 40 times) longer than tests involving IBPL.

Considerations such as speed of profile induction and classification are important. In order to induce a user profile based on observations, many examples are needed, and large log files are generated. As agent technology is applied to commercial tools such as Web browsers and e-mail filters, these issues have to be considered.

Few studies have previously been conducted which compare CN2 with instance-based algorithms that utilize the value-distance metric. PEBLS [18] is an instance-based algorithm which uses a modified version of this distance metric. Its performance has been compared to that of a number of other learning methods, such as Backpropagation and ID3 [18], C4.5 [19], CN2 [19], and naive bayesian classifiers [20] over a variety of different domains. For a general discussion of the differences between rule-based and instance-based learning systems, see [21].

8 FUTURE WORK

Our agent model is currently being applied to the task of identifying interesting information on the World Wide Web. An agent is being developed which logs pages visited by the user. From this, a user profile is induced which can be used to assist the user in two ways. As the user examines Web pages, interesting links are highlighted. This is similar to the approach used in WebWatcher [12]. In addition, a Web search engine is being developed which uses the profile to search for links to other pages of interest. These links will then be presented to the user.

Techniques such as TF-IDF [9] are being explored and compared to the feature extraction and learning techniques described above. We also hope to investigate the use of genetic algorithm techniques to develop user profiles and plan to evaluate such techniques for filtering USENet news, e-mail, and locating information on the World Wide Web.

ACKNOWLEDGMENTS

We gratefully acknowledge the authors and maintainers of *Xmail* and the authors of *xrn* for allowing us to use their software. We thank David Dyke for providing a number of amusing images, bringing UNA to life. Figs. 1, 4, 5, and 6 from *Applied Artificial In-*




Fig. 7. Narrow vs. broad classifications for *rec.humor*—CN2.




Fig. 8. Narrow vs. broad classifications for *rec.humor*—IBPL.




Fig. 9. Effect of hotlist on predictive accuracy for *sci.psychology*—CN2.

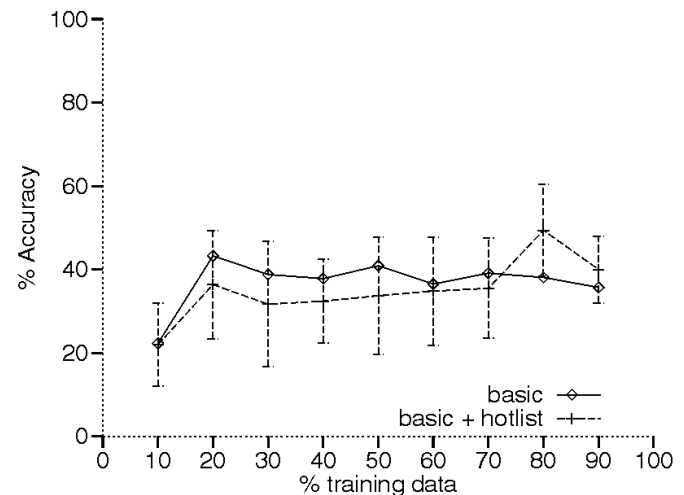


Fig. 10. Effect of hotlist on predictive accuracy for *sci.psychology*—IBPL.

telligence, vol. 11, no. 1, Terry R. Payne and Peter Edwards, Washington D.C.: Taylor and Francis Inc., reproduced with permission (all rights reserved).

Terry Payne acknowledges financial support provided by the U.K. Engineering and Physical Sciences Research Council (EPSRC).

REFERENCES

- [1] B.D. Sheth, "A Learning Approach to Personalized Information Filtering," master's thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Inst. of Technology, 1994.
- [2] T.R. Payne, "Learning Email Filtering Rules with Magi, A Mail Agent Interface," master's thesis, Dept. of Computing Science, Univ. of Aberdeen, Scotland, 1994.
- [3] T.M. Mitchell, R. Caruana, D. Freitag, J. McDermott, and D. Zabowski, "Experience with a Learning Personal Assistant," *Comm. ACM*, vol. 37, no. 7, pp. 81-91, 1994.
- [4] R. Kozierok and P. Maes, "A Learning Interface Agent for Scheduling Meetings," *Proc. ACM-SIGCHI Int'l Workshop Intelligent User Interfaces*, ACM Press, New York, pp. 81-88, 1993.
- [5] T.W. Malone, K.R. Grant, F.A. Turbak, S.A. Brobst, and M.D. Cohen, "Intelligent Information-Sharing Systems," *Comm. ACM*, vol. 30, no. 5, pp. 390-402, 1987.
- [6] D.N. Chin, "Intelligent Interfaces as Agents," *Intelligent User Interfaces*, J.W. Sullivan and S.W. Tyler, eds., ACM Press, New York, pp. 177-206, 1991.
- [7] Y. Gil, "Trainable Software Agents," *Software Agents: Papers from 1994 Spring Symp.*, Menlo Park, Calif., AAAI, pp. 99-102, 1994.
- [8] K. Lang, "NewsWeeder: Learning to Filter Netnews," *Proc. 12th Int'l Machine Learning Conference (ML95)*, Morgan Kaufmann, San Francisco, pp. 331-339, 1995.
- [9] G. Salton and M.J. McGill, *Introduction to Modern Information Retrieval*. New York: McGraw-Hill, 1983.
- [10] M.E. Metral, "Design of a Generic Learning Interface Agent," BSc thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Inst. of Technology, 1993.
- [11] C. Stanfill and D. Waltz, "Toward Memory-Based Reasoning," *Comm. ACM*, vol. 29, no. 12, pp. 1,213-1,228, 1986.
- [12] R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell, "Web-Watcher: A Learning Apprentice for the World Wide Web," *Working Notes AAAI Spring Symp. Series on Information Gathering from Distributed, Heterogeneous Environments*, Menlo Park, Calif., AAAI, 1995, <http://www.isi.edu/sims/knoblock/sss95/info-gathering.html>.
- [13] T.R. Payne and P. Edwards, "Interface Agents that Learn: An Investigation of Learning Issues in a Mail Agent Interface," *Applied Artificial Intelligence J.*, submitted for publication.
- [14] P. Clark and T. Niblett, "The CN2 Induction Algorithm," *Machine Learning*, vol. 3, pp. 261-283, 1989.
- [15] B.V. Dasarathy, *Nearest Neighbor (NN) Norms: NN Pattern Classification Techniques*. Los Alamitos, Calif.: IEEE CS Press, 1991.
- [16] C.L. Green, "USENet News Agent," BSc final year project report, Dept. of Computing Science, Univ. of Aberdeen, Scotland, 1995.
- [17] T.M. Cover, "Estimation by the Nearest Neighbor Rule," *IEEE Trans. Information Theory*, vol. 14, no. 1, pp. 50-55, 1968.
- [18] S. Cost and S. Salzberg, "A Weighted Nearest Neighbor Algorithm for Learning with Symbolic Features," *Machine Learning*, vol. 10, pp. 57-78, 1993.
- [19] P. Domingos and M. Pazzani, "Beyond Independence: Conditions for the Optimality of the Simple Bayesian Classifier," *Proc. 13th Int'l Machine Learning Conf. (ICML96)*, Morgan Kaufmann, San Francisco, pp. 105-112, 1996.
- [20] J. Rachlin, S. Kasif, S. Salzberg, and D.W. Aha, "Towards a Better Understanding of Memory-Based Reasoning Systems," *Proc. 11th Int'l Machine Learning Conf. (ML94)*, Morgan Kaufmann, San Francisco, pp. 242-250, 1994.
- [21] P. Clark, "A Comparison of Rule and Exemplar-Based Learning Systems," *Machine Learning, Meta-Reasoning, and Logics*, P.B. Brazdil and K. Konolige, eds., Kluwer, Boston, pp. 159-186, 1990.