Artificial Intelligence Review (2005) 24: 163–178
DOI 10.1007/s10462-005-4609-5

A Case-Based Explanation System for Black-Box Systems

CONOR NUGENT & PÁDRAIG CUNNINGHAM
Department of Computer Science, Trinity College, Machine Learning Group, University of Dublin, Dublin, Ireland
(E-mails:{conor.nugent, padraig.cunningham}@cs.tcd.ie)

Abstract. Most users of machine-learning products are reluctant to use them without any sense of the underlying logic that has led to the system's predictions. Unfortunately many of these systems lack any transparency in the way they operate and are deemed to be black boxes. In this paper we present a Case-Based Reasoning (CBR) solution to providing supporting explanations of black-box systems. This CBR solution has two key facets; it uses local information to assess the importance of each feature and using this, it selects the cases from the data used to build the black-box system for use in explanation. The retrieval mechanism takes advantage of the derived feature importance information to help select cases that are a better reflection of the black-box solution and thus more convincing explanations.

Keywords: black-box systems, case-based reasoning, explanation

"Computers are useless. They can only give you answers."
Pablo Picasso

1. Introduction

In machine learning research the quest for increasingly more accurate and stable classifiers has led to ever more complicated algorithms. Ensemble approaches and algorithms such as Support Vector Machines and Neural Networks have reached a level of complexity where they are not readily interpretable. Such approaches are commonly referred to as black-box algorithms owing to their lack of transparency with regard to the reasoning behind the predictions they make.

Although increases in accuracy are welcomed, recent research has highlighted the need for interpretability and transparency as a critical aspect in the implementation of machine learning techniques in real world applications (Andrews et al., 1995). People are understandably reluctant to accept without question predictions from black-box systems.

This has led to the development of explanation systems that strive to offer an insight into the workings of the black-box systems. Many different approaches have been taken but commonly the explanation systems try to build machine learning systems that are inherently interpretable such as tree-based or rule-based systems that describe the underlying black box e.g. (Andrews et al., 1995; Tickle et al., 1998; Zhou and Jiang, 2003). The relevant rules or a tree structure is then used as evidence in support of the black box's prediction. Such systems use the black box as an oracle capable of supplying an unlimited amount of training data. The hope is that, with an abundance of training data, the explanation system should offer a good description of the underlying black-box system. However, in reality such systems are limited in the level of fidelity that they can achieve while maintaining some level of interpretability. The differing bias of the black-box algorithm and that of the one being used for explanations means that it can be difficult to fully capture the operation of the black-box system. Domingos (1998) focused on how well an explanation facility captured the improvements gained through the use of ensemble techniques. He found that it retained just 60% of the gains. More accurate descriptions of the operation of the black box often come at the cost of increasingly more complex tree and rule-based systems. This trade off in interpretability versus fidelity means that such approaches are of limited use as a convincing explanation system when the underlying problem is complex and the credibility of the system can be damaged by bad, inaccurate or convoluted explanations (Majchrzak and Gasser 1991). Conversely CBR systems have an inherent transparency that has particular advantages for explanation. As Leake (1996) points out:

> "... neural network systems cannot provide explanations of their decisions and rule-based systems must explain their decisions by reference to their rules, which the user may not fully understand or accept. On the other hand, the results of CBR systems are based on actual prior cases that can be presented to the user to provide compelling support for the system's conclusions."

The use of actual training data, cases from the case base, as evidence in support of a particular prediction is a powerful and convincing form of explanation. Research by Cunningham et al. (2003) has further supported the claim that CBR explanations are more convincing than rule-based explanations in some domains. The use of a case-based explanation facility for black-box systems also helps remove the inherent fidelity/interpretability trade-off that exists in the approaches discussed previously. This has motivated us to investigate the development of a case-based explanation facility for black-box systems. This paper describes the work that we have done in developing a general framework for CBR systems and in particular an implementation for regression tasks.

The paper is structured as follows. Section 2 provides an overview of relevant work on explanation from CBR research while Section 3 outlines a general framework for CBR explanation. Section 4 introduces our case-based explanation approach for regression problems. Some examples of the proposed system in use are shown in Section 5 followed by our conclusions in Section 6.

2. Explanation in CBR and Explanations for Black Boxes

The motivation behind most explanation systems is to provide some form of evidence or argument in support of a given prediction. For instance, in a rule-based explanation system, the user will be presented with the most appropriate rule or set of rules as evidence in favour of a prediction. The success of the explanation then lies in the perceived validity of the rule presented which is not always a straightforward issue. CBR explanations provide quite a different form of explanation which has many advantages over conventional rule based approaches:

- **Use of Real Evidence:** In CBR the user is presented with actual cases, past experiences. In most applications these cases are undoubtedly true and so their validity is not in question; this is the great strength of case-based explanations. Users who are unfamiliar or suspicious of a black-box system are more likely to be convinced by explanations that contain factual evidence than by abstract rules.
- **Fixed and Simple Form of Explanation:** CBR explanations avoid the interpretability versus fidelity trade off that can plague some other techniques. The type of explanation presented to the user is independent of the complexity of the problem described by the underlying system.
- **Natural Form of Explanation:** Research in cognitive science and other areas suggests that explanation by analogy is a natural form of explanation in some domains and one that people can quickly relate to (Cunningham et al., 2003; Gentner et al., 2003).

The issue with case-based explanations lies in the perceived appropriateness of the presented cases to the validity of the prediction. This is an issue that has recently received a lot of attention in the CBR community. In CBR explanations, the ability of the user to make meaningful comparisons between the query and the retrieved explanation case is of critical importance to the success of the explanation. CBR systems are not wholly transparent and much domain knowledge can be contained within the similarity metrics used in the system. It is implicitly assumed in simple CBR explanations systems that the user has this same domain knowledge and so the appropriateness of the explanation case is clear. However, this may not be the case and the relevance of the retrieved case may be lost on novice users.

This is an issue that McSherry (2003) has addressed in his ProCon System and in an extended version of the system ProCon-2 (McSherry, 2004). McSherry has focused on making the relationship between the feature values within a case and its predicted class explicit. He argues that simply presenting the feature values in the most similar cases may be misleading. The relationship between feature values and the predicted class may not always be a positive one; the presence of some feature values may in fact be evidence against the prediction. Simply supplying the user with a case may lead them to incorrectly infer the relationship between feature values and the prediction. To combat this McSherry provides the user with extra relational information about the case feature values and the predicted class. To infer the feature-class relationships a Naïve Bayes model is built on the entire training set and from this the relational information is derived. Although the provision of extra feature relation evidence helps make the relevance of retrieved explanation cases clear the manner in which the evidence is derived has to be carefully considered. The use of global models such as Naïve Bayes to describe the CBR system may not be always appropriate as it risks the same fidelity problems seen in other explanation systems as discussed in Section 1. As Sørmo and Cassens (2004) point out, ProCon may fail to capture the effects of interactions between features in some domains.

In other work, Doyle et al. (2004) have focused on the observation that the nearest retrieved case in a CBR system may not be the best case to present as an explanation. They argue that in classification tasks cases that are between the query case and the decision boundary provide more convincing explanations. That is, cases that are more marginal on the important criteria are more convincing. With such cases the user is better able to assess whether the classification of the target case is justified.

The work of Doyle et al. (2004) and McSherry (2003) has highlighted important issues relating to case-based explanations as well as proposing solutions. It can be seen that providing the appropriate cases as well as highlighting the salient feature-class relationships within those cases are important factors in designing a successful case-based explanation facility.

3. A General Framework for CBR Explanations of Black Box Systems

As was discussed in Section 2 the provision of additional information describing the relationship between case feature-values and the prediction has the potential to greatly improve CBR explanations and address some of their shortcomings. This is particularly important when considering CBR explanations for black-box systems. Many of these systems are used because of their ability to accurately model non-linear problems. The non-linear nature of the underlying problems may mean that the relationship between features and the prediction may vary across the feature space. Some features may be important in some areas of the feature space and not at all relevant in other areas. We would like to provide the user with feature salience information that informs the user of what feature values are important in different individual predictions. It would be useful to rank each feature based on the impact it had on a given prediction and whether that impact was negative or positive. This would provide the user with a sense of the relationship of the feature values to the prediction for the presented case that they can then critically assess. These rankings will also focus the user's attention on the more important features of a case.

However, in the Blood Alcohol Content and Bronchiolitis domains that we have considered (Cunningham et al., 2003; Walsh et al., 2004) it seems that simple global models are not suitable for such tasks. To overcome this problem we have designed a two stage explanation system that takes advantage of our ability to use the black box as an oracle from which we can extract information about its behaviour in particular regions of the feature space. This information is then used to inform the user of the salience of the various features and also to drive the retrieval process. We will now discuss each of the two stages in turn.

3.1. *Derive local feature information*

The provision of feature rankings provides the user with a sense of how each of the feature values contributed to the particular prediction. It is important too that these rankings should reflect the locality of the presented case with which the prediction is made. In order to provide such feature rankings, two distinct steps are taken. Firstly the black box is treated as an oracle and an artificial data set is constructed around the point of inquiry and secondly a model is built on this data.

The black box allows us to get a prediction for any set of feature values we care to imagine. We can present the black box with feature value sets similar to those of the query case and so we can build up a local case base around the original query point. This is done by perturbing, in a controlled manner, the feature values of the case we are providing an explanation for and using the black box to attach a prediction to the artificial case. In the case of the regression implementation outlined in Section 4 for each artificial case to be created a single feature from the query case is chosen and its value is altered. This is done by adding a controlled amount of noise to its value if it is continuous or replacing it with an alternative value if a nominal feature. This new artificial case is then classified using the black box. This process is repeated a number of times for each feature until there is enough data for an accurate model of the black box's function in that area of the feature space to be built. The number of artificial cases needed to be created is determined experimentally.

The choice of model used to capture the local behaviour is primarily driven by the need for a model that can provide us with information about each feature's relative importance and also its relationship to the predicted value. We would also like the chosen model to accurately describe the black box in the local region being investigated, although the use of a local artificial case base goes a long way towards achieving this task. The feature salience information derived from the chosen model can then be used to help select the candidate explanation case from the original data used to build the black-box system.

3.2. *Case retrieval mechanism*

Ideally we would like to present the user with cases that reinforce the black box's prediction. However, the non-linear nature of the relationship of features to the prediction also has implications for the selection of cases to present to the user. The feature rankings may indicate that some features are more important than others and this should be reflected in the retrieval process.

One simple way to utilise this information would be by simply weighting each of the features based on their relative importance. For instance, imagine a simple two feature problem has been learnt by a black box and we would like to select a case to use as an explanation of a prediction given for a particular set of inputs, QP. As can be seen in Figure 1a, if the features are un-weighted, C1 is the nearest neighbour.

However imagine from our feature ranking information we discover that feature two is more important than feature one. The rankings mean that from the black box's perspective, feature two has a greater impact on the predicted value than feature one. This means that cases that are closer in value to QP in feature two bear greater relation to it and so these are the cases we should seek out. By warping the axis using the feature weights as in Figure 1b, greater emphasis can be put on this feature and a different nearest neighbour, C2, is found.

3.3. *Outline of the general explanation process*

The flow of execution and the relationships and dependencies between individual processes in the framework outlined above is described in Figure 2. In our explanation framework each explanation is tailored to the particular set of inputs and the prediction made. First the Query Case is used to seed the generation of an artificial case base.

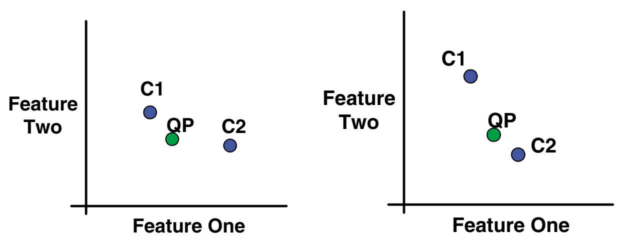

(a) Un-Weighted Axis **(b)** WeightedAxis

Figure 1. QP and its neighbours with weighted and un-weighted axis.

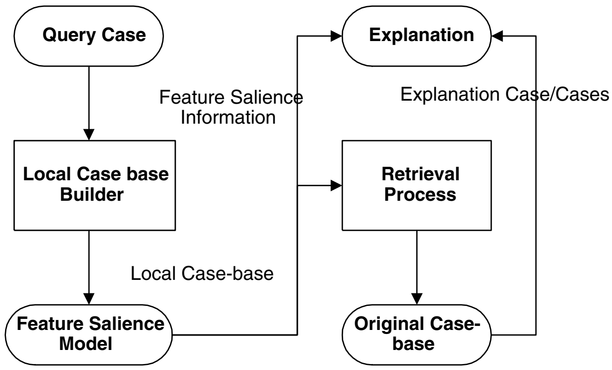

Figure 2. A flow diagram of the explanation process.

Once we have this data we then describe it using some model as discussed in Section 3.1. We now have a description of the behaviour of the black-box model in the region of interest and this information is then used in two further stages of the explanation process, the case retrieval and the explanation stage, as can be seen below. In the case retrieval process the feature salience information is used to select the best case from the original case base to use in the explanation. This is then passed on to the final explanation stage where it, along with the feature salience information, is used to generate the final explanation presented to the user. The exact form of the explanation and what information is presented to the user is very much both task and domain dependent (Cassens, 2004; Sørmo et al., 2005). In Section 5 we present one possibility in which the feature salience information and retrieved case are used to generate a discursive text justifying a black box's prediction. We believe that the provision of CBR explanations for black-box systems based on the local derived feature ranking and the presentation of appropriately selected cases as discussed will provide users of black-box systems with satisfactory explanations. In the following section we discuss the implementation of such a system for regression tasks.

4. An Explanation System for Regression

Although this paper advocates case-based explanations for black-box systems for both classification and regression, the following discusses the implementation for a regression system. As discussed above there

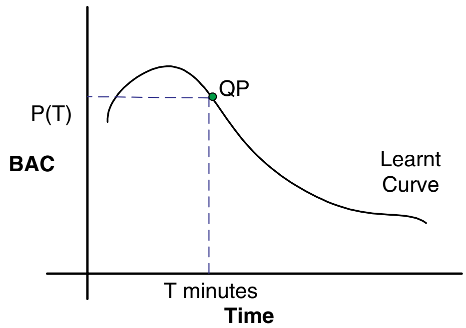

Figure 3. The function learnt by the NN-BAC vs. time.

are two important tasks that are integral to the explanation system, local feature salience information and the provision of cases that are appropriate given the feature rankings. Previously these tasks were discussed in an abstract sense but we will now discuss in concrete terms how these objectives can be achieved in regression problems.

4.1. *Local feature ranking*

As an example of how this might be done imagine we have a neural network model that predicts the Blood-Alcohol content (BAC) in a person's blood after they have consumed a certain number of units of alcohol and stopped drinking. The graph of the function learnt by the neural network (NN) might look something like the one in Figure 3. As the consumed units are absorbed into the body the BAC value increases until it has reached a maximum value from where the level then begins to fall back down as the body processes the alcohol. The function learnt by the NN is of course unknown to us and so when we ask it to provide a prediction for time T we will simply be presented with a prediction P(T) with no insight on how this prediction was derived. We can then begin to proposition the NN with cases similar to our query case (QP) and build a case base that describes the NN's function around QP as seen in Figure 4.

Once we have built an artificial case-base around QP that accurately describes the black box's function in that area we are then left with the problem of how best to extract feature rankings from it. For regression tasks, multivariate linear regression models would seem to be the best candidate for deriving such information. A linear regression model provides us with a set of coefficients for each feature that

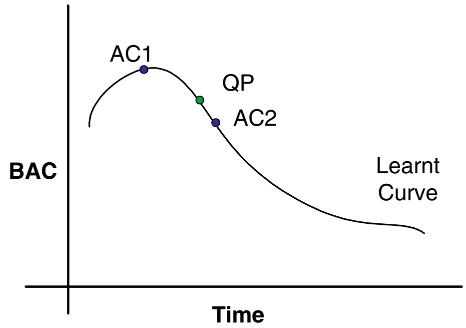

Figure 4. Artificial data points AC1 and AC2 are created around QC.

can then be used to infer how sensitive the prediction is to changes in each feature's value and so its relative importance. The coefficients also provide information about whether a feature is negatively or positively correlated with the prediction variable at that point. In our particular example the coefficient would give us the rate at which BAC is changing with time at that particular point. However, care must be taken to ensure the linear model derived truly reflects the NN's function. If we were simply to build our model on the locally built case base without attention to each case's relation to a query case we would end up with a model like that shown in Figure 5. This would be an un-weighted linear model and is not a good model of the NN's behaviour at point QP. To overcome such problems locally weighted linear regression can be used (Atkeson et al., 1997). Local linear regression allows us to weight each case based on its similarity to the query case. In the case of our implementation we use an un-weighted Euclidean distance measure as the weighting function. For instance AC1 would be given a lower weight than AC2 and so would

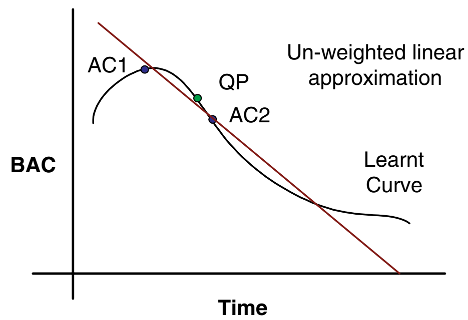

Figure 5. Fitting a linear model to the artificially created data.

have less of an impact on the derived model. This gives us a model that is close to a tangent to the curve at QP and gives us a slope value that truly reflects the NN's function as can be seen in Figure 6. The above example is quite simple and the information extracted may not seem to be very useful, but in a multi-dimensional problem such information is extremely useful. In such a case, a hyperplane is produced and each coefficient of that model gives us a sense of how each feature relates to the predicted value. How the feature salience information is used in the final explanation stage is very much dependent on the context the system is being used in, on what is deemed most effective and useful for a particular domain, and the users in that domain. In Section 5 we give one possible example of how this information can be used.

4.2. *The provision of appropriate cases*

As has been stated before the strength of CBR explanations lies in the use of previous experience, of actual training cases. However it is important to provide the user with the appropriate cases that support the prediction. Once a set of feature rankings has been derived it is quite a simple task to adjust the selection of cases. A nearest neighbour algorithm is used to select cases from the original training data used to build the NN. Each feature is weighted based on the magnitude of the coefficient given to it and the nearest neighbour algorithm is then applied using these weights. This process helps eliminate the noise introduced by features that are not relevant to the particular case for which we are providing an explanation.

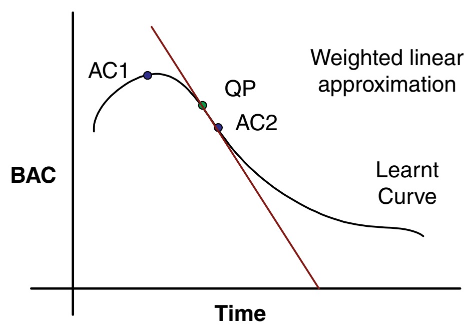

Figure 6. Fitting a locally weighted linear model.

5. Sample Explanations

As an example of how the salience information derived using our regression explanation scheme can be used to provide convincing explanations we applied it to the task of explaining the predictions of a NN. The NN was trained to predict BAC. The task involves using information about people's weight, gender, number of units of alcohol consumed, etc. to predict the concentration of alcohol in their blood stream. The training data was taken from the data that had previously been collected and used by Cunningham et al. (2003). The full set of features used can be seen in Table 1. Extra data was added in order to ensure that the NN learnt the underlying function over the full feature space. The extra data added were examples of BAC when no units were consumed. This ensured that the NN learnt effectively to deal with the special but important scenario of when no units have been consumed.

Table 1. The features in the BAC data set

Weight (Kg)	Duration (Time spent drinking)
Meal (None, Snack, Lunch, Full)	Amount (In units)
Gender	BAC (Blood alcohol content)

For this particular task we felt that the most effective of explanation would be one in the form of a discursive text. The text is derived from the feature salience information and contains two sections:

- **Provision of feature importance information.** In this first section we inform the user of what the particularly important features are. The feature importance information is derived from the absolute values of the coefficients of the local linear model. These values reflect the sensitivity of the predicted value to changes in those features. Features which the predicted value is particularly sensitive to can be deemed important and we form a ranked list of such features. If the coefficients are particularly small and below a fixed threshold then that feature is determined not to be important to that particular prediction. We then present these to the user using a simple template.
- **Explanation of feature-value differences.** As McSherry (2003) has observed it is important to inform the users of the effects of

differences. Using the local linear model it is straightforward task to determine the effects of feature value differences. The feature value differences can be inserted into the model and their effects observed as the changes in the model output. Feature differences which have very little or no effect on the BAC value are ignored. A simple template is used to present this information to the user.

The discursive text produced can be seen in Tables 2–4. In each case we would deem the retrieved case and the discursive text to be a convincing explanation. It is clear that the feature information is in line with our intuitive understanding of the problem and that they add value to the overall explanation. The text is useful in appreciating the difference between the predicted BAC value and that in the explanation case as well as offering an insight into the nature of the problem being studied.

In Example A we can see that the Amount and Duration features are the two most important factors in this prediction. Although the retrieved explanation case is a reasonably close match there are some differences in the feature values as well as a difference in the predicted value for the Query Case and the BAC value of the explanation case. These differences are accounted for in the second section of the discursive text. The user is informed that both Duration being bigger and Amount being smaller in the Explanation Case have the effect of decreasing the BAC value. These differences then account for the

Table 2. Example A

	Query case	Explanation case
Weight(kgs)	79.0	70.0
Duration (mins)	200.0	270.0
Gender	Male	Male
Meal	Lunch	Lunch
Amount (Units)	9.9	9.6
BAC	26 (predicted)	21

Explanation:
The important features in determining this prediction, listed in order of impact, are: Amount, Duration, Gender, Weight and Meal.

Duration being bigger in the explantaion case has the effect of decreasing its BAC. Amount being smaller in the explantaion case has the effect of decreasing its BAC value.

Table 3. Example B

	Query case	Explanation case
Weight (kgs)	85	80
Duration (mins)	60	60
Gender	Male	Male
Meal	Lunch	Lunch
Amount (Units)	4.9	4.9
BAC	12 (predicted)	14

Explanation:
The important features in determining this preidction, listed in order of impact, are: Amount, Duration, Gender, Weight and Meal.

Weight being smaller in the explanation case has the effect of increasing its BAC value.

Table 4. Example C

	Query case	Explanation case
Weight (kgs)	73	69
Duration (mins)	210	200
Gender	Male	Male
Meal	Lunch	Lunch
Amount (Units)	0	0
BAC	0 (predicted)	0

Explanation:
The important features in determining this prediction, listed in order of impact, are: Amount.

There were no significant differences between the Explanation and Query case.

smaller BAC value of the explanation case. Although there is a difference in the weight of the two subjects this difference did not contribute significantly to the difference in BAC value.

Again in Example B we can see that the discursive text produced is useful. Here the Query Case and the Explanation Case are almost identical except for a difference in weight. In this case the difference in weight is significant and accounts for the Explanation Case BAC value being slightly higher.

Example C highlights the value of deriving feature salience locally. There is one notable discontinuity in the BAC dataset at the point when the number of units consumed is zero. Clearly at this point an individual's weight or gender etc. will have no altering effect on their BAC level. This is an extreme example of the effect interactions and dependencies between features can have. In Example C we can see that by deriving our feature salience information locally we can provide an explanation that captures this intuitive fact. Amount is found to be the only significant feature in this area of the feature space. None of the differences in the Explanation Case and the Query Case were found to have any effect in changing the BAC value and so there were no significant differences.

6. Conclusion

Providing useful explanations for black-box systems is an important issue and one for which we feel CBR is ideally suited. We have highlighted the important issues involved in the application of CBR explanations to black boxes as well as outlining possible solutions to these problems. In particular we focused on an implementation of a CBR explanation system for regression tasks. We believe the explanations produced through this system to be straightforward, useful, and convincing, avoiding many of the pitfalls that can plague other approaches. Our approach is independent of the particular black-box model being used and requires only that the training data used to build the model is retained. This has encouraged us to further investigate the use of CBR for explanations of black-box systems.

In the future we will expand our methods to classification problems using logistic regression as a local model from which to derive feature salience information and a measure of confidence in a given prediction. We would also like to focus on improved methods of case retrieval and of generating local artificial data.

References

Andrews, R., Diederich, J. & Tickle, A. (1995). A Survey and Critique of Techniques for Extracting Rules from Trained Artificial Neural Networks. *Knowledge Based Systems* **8**: 373–389.

Atkeson, C. G., Moore, A. W. & Schaal, S. (1997). Locally Weighted Learning. *Artificial Intelligence Review* **11**: 11–73.

Cassens, J. (2004). Knowing What to Explain and When. In Gervás, P. & Gupta, K. (eds.) *Proceedings of the ECCBR 2004 Workshops*. 97–104, Technical Report 142-04, Departmento de Sistemas Informáticos y Programación, Universidad Complutense de Madrid, Madrid, Spain.

Cunningham, P., Doyle, D. & Loughrey, J. (2003). An Evaluation of the Usefulness of Case-Based Explanation. In Ashley, K. D. & Bridge, D. G. (eds.) *ICCBR*, Vol. 2689 of *Lecture Notes in Computer Science*. 122–130, Springer: Berlin.

Domingos, P. (1998). Knowledge Discovery Via Multiple Models. *Intelligent Data Analysis* **2**(1–4): 187–202.

Doyle, D., Cunningham, P., Bridge, D. & Rahman, Y. (2004). Explanation Oriented Retrieval. In Funk, P. & Calero, P. (eds.) *Advances in Case-Based Reasoning (Procs. of the Seventh European Conference on Case-Based Reasoning)*. 157–168, Springer: Berlin.

Gentner, D., Loewenstein, J. & Thompson, L. (2003). Learning and Transfer: A General Role for Analogical Encoding. *Journal of Educational Psychology* **95**(2): 393–408.

Leake, D. (1996). CBR in Context: The Present and Future. In Leake, D. (ed.) *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. 3–30. AAAI Press, Menlo Park, CA.

Majchrzak, A. & Gasser, L. (1991). On using Artificial Intelligence to Integrate the Design of Organizational and Process Change in US Manufacturing. *AI and Society* **5**: 321–338.

McSherry, D. (2003). Explanation in Case-based Reasoning: an Evidential Approach. In *Proceeding of 8th UK Workshop on Case-Based Reasoning*. 47–55.

McSherry, D. (2004). Explaining the Pros and Cons of Conclusions in CBR. In Funk, P. & Calero, P. (eds.) *Advances in Case-Based Reasoning (Proceedings of the Seventh European Conference on Case-Based Reasoning)*. 317–330, Springer: Berlin.

Sørmo, F. & Cassens, J. (2004). Explanation Goals in Case-Based Reasoning. In Gervás, P. & Gupta, K. (eds.) *Proceedings of the ECCBR 2004 Workshops*. 165–174, Technical Report 142-04, Departmento de Sistemas Informáticos y Programación, Universidad Complutense de Madrid, Madrid, Spain.

Sørmo, F., Cassens, J. & Aamodt, A. (2005). Explanation in Case-Based Reasoning: Perspectives and Goals. *Artificial Intelligence Review*. This Issue.

Tickle, A., Andrews, R., Golea, R. & Diederich, J. (1998). The Truth Will Come to Light: Directions and Challenges in Extracting Rules from Trained Neural Networks. *IEEE Transactions on Neural Networks* **9** 1057–1068.

Walsh, P., Cunningham, P., Rothenberg, S. J., O'Doherty, S., Hoey, H. & Healy R. (2004). An Artificial Neural Network Ensemble to Predict Disposition and Length of Stay in Children Presenting with Bronchiolitis. *European Journal of Emergency Medicine* **11**(5): 259–264.

Zhou, Z.-H. & Jiang Y. (2003). Diagnosis with C4.5 Preceded by Artificial Neural Network Ensemble. *Transactions on Information Technology in Biomedicine* **7**(1): 37–42.