Efficient and Effective Case Reject-Accept Filtering: A Study Using Machine Learning

Robert BEVAN, Alessandro TORRISI, Katie ATKINSON, Danushka BOLLEGALA
and Frans COENEN

Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, UK
e-mail: {robert.bevan, alessandro.torrisi, k.m.atkinson, danushka.bollegala,
coenen}@liverpool.ac.uk

Abstract. The decision whether to accept or reject a new case is a well established task undertaken in legal work. This task frequently necessitates domain knowledge and is consequently resource expensive. In this paper it is proposed that early rejection/acceptance of at least a proportion of new cases can be effectively achieved without requiring significant human intervention. The paper proposes, and evaluates, five different AI techniques whereby early case reject-accept can be achieved. The results suggest it is possible for at least a proportion of cases to be processed in this way.

Keywords. Early Case Accept-Reject, Heuristics, Machine Learning

1. Introduction

A common challenge for any commercial law firm is deciding whether to take on a case or not. There are many issues of significance here; however, in most cases the primary concerns are the resource required to process a case and the potential economic gains that may result. To arrive at a decision, lawyers use their experience accompanied by established heuristics. The nature of these heuristics is frequently dependent on the nature of the area in which a particular legal concern operates. One example is to reject cases where some statutory limitation is imminent, another is to reject cases which have an international aspect. Whatever the case, the established approach is time consuming, subjective and requires reference to domain experts. However, most legal enterprises have substantial repositories of previous cases. These repositories can be effectively used to determine the potential outcome of a case. There has been a substantial amount of previous work directed at the use of machine learning to build classification models that can be used to predict the legal outcome of new cases [1,2,3,4,5,6]. However, these approaches either require substantial analysis of the new case so as to identify the key features, or operate on highly structured text composed by domain experts.

In this paper it is suggested that a more efficient approach is to use a classification model founded on features extracted from information that can be readily ascertained by an administrator on first contact with a potential client. In this way, an early "definite

reject" or "definite accept" decision can be made with respect to at least a proportion of cases, with the remaining cases set aside for further consideration. It is conjectured that this approach will be of particular benefit in the context of legal firms that operate in well defined domains, such as road traffic or household insurance litigation, where many of the cases received fall into a limited number of categories. More specifically, this paper considers five different techniques whereby this early reject/accept decision can be implemented. The first technique considered is a simple rule-based approach, using established heuristics, to filter incoming cases into three classes: *definite reject*, *definite accept* and *further consideration*. The second is to use an existing case repository to build a classification model, directed at the above three classes, using established machine learning algorithms [7]. The third is founded on the idea of first clustering the cases held in the repository into a small number of clusters and then building individual classification models for each cluster. The fourth is then a combination of the first and second approaches, whilst the fifth is a combination of the first and third.

2. Rule-Based Early Reject-Accept

The rule-based approach is the simplest and is founded on the observation that lawyers frequently apply heuristics to decide whether to accept or reject a case. These heuristics are designed to capture potential risks/rewards associated with a case. For example, if a case has an international element, pursuing the case may incur additional costs, reducing the potential reward. Conversely, there may exist some domain-specific green flags that indicate a high chance of success.

The business benefits derived from each heuristic depend on how frequently they are satisfied and how strongly they indicate whether a case should be accepted or rejected. With this in mind, two heuristics were chosen with respect to the experiments presented in Section 6: reject any cases with an international element, and reject any cases in which a statutory limitation will be reached within a specified time period. The international element heuristic was deemed to be satisfied if the initial contact indicated an overseas context. The limitation heuristic was implemented as follows: every date mentioned in the claimant's interview statement was extracted and ordered chronologically and the most recent date used to provide a conservative estimate of whether some limitation was imminent.

3. Single Classifier Based Early Reject-Accept

The idea underpinning the single classifier-based technique is to use the case repository, that most legal firms maintain, to build a classification model. In this work, classification features included in a telephone interview were converted into a bag-of-words representation, composed of unigrams and person/organisation entities. A large number of algorithms can be used to build such classification models. The authors experimented with a number of these and found there was little to choose between them. For the evaluation presented in Section 6, a regularized logistic regression model was used.

4. Multiple Classifier Based Early Reject-Accept

The idea underpinning the multiple classifier-based technique was that many of the cases typically received by legal firms can be categorised as being of a certain type and this categorisation is frequently an indicator of whether a case should be accepted or rejected. With respect to the third technique, the idea was thus to first cluster the cases in a given repository into a small number of clusters and to then construct individual classification models with respect to the cases in each cluster. The intuition was that this might produce more accurate predictions.

In order to effectively cluster the dataset, cases were initially converted into topic vectors, with the topics constructed in a semi-supervised fashion [8]. The topic-word distributions were then inspected, and all but the four most coherent topics were discarded (ie. those that best represented particular case types). The reduced topic representations were then partitioned into five clusters using the k-means algorithm [9]. After clustering, cluster-specific classification models were trained using the same features as described in Section 3. Again, a regularized logistic regression model was used for the evaluation presented below.

5. Evaluation Data

The dataset used for experimentation comprised 40,000 accident claim records collected over a period of four years. Each record contained two parts: a free-text statement transcribed from a telephone interview, and a second, shorter text summarising the claim. Each record was accompanied by a label indicating whether or not the case was accepted. Note that the dataset classes were imbalanced, with the majority of cases rejected. The first technique, the rule-based technique, was applied to the entire dataset. For the classification based techniques, techniques two and three, both training and test sets were required. Thus evaluation results in these two cases were generated using Ten-Fold Cross-Validation. As the labels included in the dataset indicate only whether a case was accepted or rejected, some extra work was required in order to classify cases into the *definite reject*, *definite accept* and *further consideration* classes. For this purpose the test set class membership probabilities were sorted, and the top 10% most confident predictions for the accept and reject classes were selected as the *definite accept* and *definite reject* class predictions, the remaining 80% were assigned to the *further consideration* class. For the combination techniques, techniques four and five, heuristic filtering was applied to the test set in addition to invoking a classification model, and the heuristic filtering class predictions took precedence over the classification model predictions.

6. Results

The results obtained from the comparative evaluation are given in Table 1. The results presented are with respect to the *definite reject* and *definite accept* classes. From the table it can be seen that the rule-based approach achieved the highest precision but also produced the lowest recall by a large margin. This method fails to make any *definite accept* predictions as the chosen heuristics target the *definite reject* class only. Both ma-

chine learning approaches increased the recall at the expense of precision. No significant performance difference was observed between the general classifier and pre-clustering approaches. This may be a result of the clustering approach assigning the majority of examples to a single broad cluster, meaning the majority of test examples were evaluated using this cluster's classification model, which is nearly equivalent to the general classification model. Pre-filtering cases using the rule-based method before applying either machine learning approach improved performance. Note the disparity in performance across classes can be partially explained by the class imbalance in the dataset.

Table 1. Comparative Evaluation Results. Metrics computed separately for the *definite accept* (+) and *definite reject* (-) classes. Note the final column (MCC) refers to the Mathews correlation coefficient.

Technique	Prec (+)	Prec (-)	Rec (+)	Rec (-)	F1 (+)	F1 (-)	MCC
1. Rule-based	-	**0.946**	-	0.041	-	0.078	**-**
2. Single Classifier-Based	0.445	0.936	**0.875**	0.628	**0.59**	0.752	0.438
3. Multiple Classifier-Based	0.439	0.934	0.869	0.625	0.584	0.749	0.42
4. Rule & Single Classifier-Based	**0.449**	0.937	0.844	**0.693**	0.586	**0.797**	0.455
5. Rule & Multiple Classifier-Based	0.445	0.934	0.836	0.691	0.581	0.794	**0.457**

7. Conclusions

In this paper five alternative techniques have been proposed directed at the task, frequently undertaken by legal firms, of deciding whether to accept or reject a new case. The techniques were evaluated using an accident claims dataset of 40,000 records. The evaluation indicated it is easily possible to implement a case reject-accept system that can automatically process at least a proportion of cases. In practice, the volume of cases an automated early-rejection/acceptance system is able to process will depend on the domain and business strategy of the operator. For future work the authors intend to investigate a confidence scoring mechanism that can be attached to a particular classification and an explanation facility.

References

[1] D. Arditi and T. Pulket, "Predicting the outcome of construction litigation using an integrated artificial intelligence model," *Journal of Computing in Civil Engineering*, vol. 24, no. 1, pp. 73–80, 2009.

[2] S. Brüninghaus and K. D. Ashley, "Combining case-based and model-based reasoning for predicting the outcome of legal cases," in *Proc. International Conference on Case-Based Reasoning (ICCBR '03)*, pp. 65–79, Springer LNAI, 2003.

[3] K. Ashley and S. Brüninghaus, "Automatically classifying case texts and predicting outcomes," vol. 17, pp. 125–165, 06 2009.

[4] M. Dunn, L. Sagun, H. Şirin, and D. Chen, "Early predictability of asylum court decisions," in *Proceedings of the 16th Edition of the International Conference on Artificial Intelligence and Law*, ICAIL '17, (New York, NY, USA), pp. 233–236, ACM, 2017.

[5] D. L. Chen and J. Eagel, "Can machine learning help predict the outcome of asylum adjudications?," in *Proceedings of the 16th Edition of the International Conference on Artificial Intelligence and Law*, ICAIL '17, (New York, NY, USA), pp. 237–240, ACM, 2017.

[6] N. Aletras, D. Tsarapatsanis, D. Preotiuc-Pietro, and V. Lampos, "Predicting judicial decisions of the european court of human rights: a natural language processing perspective," *PeerJ Computer Science*, vol. 2, p. e93, 2016.

[7] M. Bramer, *Principles of Data Mining*. Springer, 3 ed., 2016.

[8] Jagadeesh Jagarlamudi, Hal Daumé III, and Raghavendra Udupa, "Incorporating lexical priors into topic models," pp. 204–213, Association for Computational Linguistics, 2012.

[9] J. B. MacQueen, "Some methods for classification and analysis of multivariate observations," p. 281297, University of California Press, 1967.