

Set-Valued Prediction in Multi-Class Classification

Thomas Mortier¹, Marek Wydmuch², Krzysztof Dembczyński²,
Eyke Hüllermeier³, and Willem Waegeman¹

¹ Department of Data Analysis and Mathematical Modelling,
Ghent University, Belgium

{thomasf.mortier, willem.waegeman}@ugent.be

² Institute of Computing Science,
Poznań University of Technology, Poland
marek@wydmuch.poznan.pl, kdembczynski@cs.put.poznan.pl

³ Intelligent Systems and Machine Learning,
Universität Paderborn, Germany
eyke@upb.de

Abstract. In cases of uncertainty, a multi-class classifier preferably returns a set of candidate classes instead of predicting a single class label with little guarantee. More precisely, the classifier should strive for an optimal balance between the correctness (the true class is among the candidates) and the precision (the candidates are not too many) of its prediction. We formalize this problem within a general decision-theoretic framework that unifies most of the existing work in this area. In this framework, uncertainty is quantified in terms of conditional class probabilities, and the quality of a predicted set is measured in terms of a utility function. We then address the problem of finding the Bayes-optimal prediction, i.e., the subset of class labels with highest expected utility.

Keywords: Multi-class classification · Set-valued prediction · Expected utility maximization.

1 Introduction

In probabilistic multi-class classification, one often encounters situations in which the classifier is uncertain about the class label for a given instance. In such cases, instead of predicting a single class, it might be beneficial to return a set of classes as a prediction, with the idea that the correct class should at least be contained in that set. For example, in medical diagnosis, when not being sure enough about the true disease of a patient, it is better to return a set of candidate diseases. Provided this set is sufficiently small compared to the total number of diagnoses, it can still be of great help for a medical doctor, because only the remaining candidate diseases need further investigation.

Multi-class classifiers that return set-valued predictions have been considered by several authors under different names. In [3], the notion of *non-deterministic classification* is introduced, and the performance of set-valued classifiers is assessed using set-based utility scores from the information retrieval community, such as precision, recall, and the F_1 -measure. Unlike typical top- t prediction in information retrieval, the cardinality t of the set varies over instances in set-valued prediction, depending on the uncertainty about the class label for an instance. Other researchers call the same setting *credal* or *cautious classification*. In a series of papers, they analyze several set-based utility scores that reward abstention in cases of uncertainty [1, 2, 8, 7]. The framework of *conformal prediction* also produces set-valued predictions, albeit with a focus on confidence (the set covers the true class with high probability) and less on utility [6]. Furthermore, set-valued prediction can be seen as a generalization of multi-class classification with a reject option [5], where one either predicts a single class or the complete set of classes.

2 Proposed framework

As main contribution[4], we develop efficient Bayes-optimal algorithms for maximizing a broad family of set-based utility scores. To this end, a decision-theoretic framework is proposed, consisting of estimating a conditional class distribution, which can be solved by means of several types of probabilistic methods, such as neural network-based softmax classifiers. At prediction time, we then use the estimated probabilistic model in a subsequent inference procedure, by means of expected utility maximization. Due to an exponential time complexity, as we need to consider all possible subsets of classes, the inference procedure results in a non-trivial optimization problem. However, in our work, we show that this problem can be solved in an efficient manner. In fact, by making a restriction to sets of classes that occur as nodes in a predefined hierarchy, we are able to further improve the efficiency of the inference procedure. By analyzing the number of queries to the probabilistic model, we show theoretically and empirically that gains in runtime (often, but not always) have the price of losses in predictive performance.

3 Conclusion

We briefly discussed set-valued prediction in multi-class classification by means of a decision-theoretic framework, for a broad family of set-based utility functions. In this framework, the Bayes-optimal solution is found by using expected utility maximization. Although resulting in a non-trivial optimization problem, we argue that this can be solved in an efficient manner. By further exploiting search space restrictions, we propose other Bayes-optimal inference algorithms that improve runtime efficiency, however, keeping in mind the price that one has to pay in predictive performance.

References

1. Corani, G., Zaffalon, M.: Learning reliable classifiers from small or incomplete data sets: The naive credal classifier 2. *Journal of Machine Learning Research* **9**, 581–621 (2008)
2. Corani, G., Zaffalon, M.: Lazy naive credal classifier. In: *Proceedings of the 1st ACM SIGKDD Workshop on Knowledge Discovery from Uncertain Data*. pp. 30–37. ACM (2009)
3. Coz, J.J.D., Díez, J., Bahamonde, A.: Learning nondeterministic classifiers. *The Journal of Machine Learning Research* **10**, 2273–2293 (2009)
4. Mortier, T., Wydmuch, M., Dembczyński, K., Hüllermeier, E., Waegeman, W.: Efficient algorithms for set-valued prediction in multi-class classification (June 2019), arXiv:1906.08129
5. Ramaswamy, H.G., Tewari, A., Agarwal, S.: Consistent algorithms for multiclass classification with a reject option. *CoRR* **abs/1505.04137** (2015)
6. Shafer, G., Vovk, V.: A tutorial on conformal prediction. *Journal of Machine Learning Research* **9**, 371–421 (2008)
7. Yang, G., Destercke, S., Masson, M.H.: The costs of indeterminacy: How to determine them? *IEEE Transactions on Cybernetics* **47**, 4316–4327 (2017)
8. Zaffalon, M., Giorgio, C., Mauá, D.D.: Evaluating credal classifiers by utility-discounted predictive accuracy. *Int. J. Approx. Reasoning* **53**, 1282–1301 (2012)