

Improving the Decision Support in Diagnostic Systems using Classifier Probability Calibration

Xiaowei Kortum¹, Lorenz Grigull², Urs Muecke², Werner Lechner³, and Frank Klawonn^{1,4}

¹ Department of Computer Science,
Ostfalia University of Applied Sciences
Salzdahlumer Str. 46/48, 38302 Wolfenbuettel, Germany
{x.kortum, f.klawonn}@ostfalia.de

² Department of Paediatric Haematology and Oncology,
Medical University Hanover
Carl-Neuberg Str.1, 30625 Hannover, Germany
{grigull.lorenz, urs.muecke}@mh-hannover.de

³ Improved Medical Diagnostics IMD GmbH,
Ostfeldstr. 25, Hannover 30559, Germany
werner.lechner@improvedmedicaldiagnostics.com

⁴ Helmholtz Center for Infection Research
Inhoffenstrasse 7, 38124 Braunschweig, Germany
Frank.Klawonn@helmholtz-hzi.de

Abstract. In modern medical diagnoses, classifying a patient's disease is often realized with the help of a system-aided symptoms interpreter. Most of these systems rely on supervised learning algorithms, which can statistically extend the doctor's logic capabilities for interpreting and examining symptoms, thus supporting the doctor to find the correct diagnosis. Besides, these algorithms compute classifier scores and class labels that are used to statistically characterize the system's confidence level on a patient's type of disease. Unfortunately, most classifier scores are based on an arbitrary scale. Thus the interpretations often lack clinical significance. Especially when combining multiple classifier scores within a diagnostic system, it is essential also to apply a calibration process to make the different scores comparable.

As a frequently used calibration technique, we adapted isotonic regression for our medical diagnostic support system, to provide a flexible and effective scaling process that consequently calibrates the arbitrary scales of probability scores from classifiers. In a comparative evaluation, we show that isotonic regression can actively improve the system diagnosis based on an ensemble of classification results, also effectively remove outliers from data, thus optimize the decision support system to obtain more accurate diagnostic results.

Keywords: Classifier Calibration, Isotonic Regression, Pool Adjacent Violators, Multiple-Classifer System, Statistical Computing

1 Introduction

Compared with the common disease appearances, diagnosing rare types based on a patient’s symptoms is hard to achieve. In particular, when multiple influencing factors need to be taken into account, overlooked or inaccurately interpreted symptoms of a rare disease are common. Even when the preliminary examination for a patient’s physical condition and laboratory tests have been completed for delineating the range of possible diseases, potential lack of awareness and data resources make rare diseases still hard to be identified for many medical doctors. De facto, classifying rare diseases involves a sophisticated process which can take years from early symptom appearance to a final diagnosis. Solutions like a system-assisted diagnosis for improved interpretation of signs through collective patient records are a frequently requested strategy by medical institutions.

A collaborative research project between scientific researchers and medical experts from Hannover Medical School (MHH) has been initiated. The original idea of this cooperation based on realizing a computer-aided diagnostic system that supports medical doctors’ classification capabilities on new patients with symptoms for specific types of rare disease. For effectively collecting patients data, MHH designed questionnaires for particular disease groups based on interviews, investigations, and observations of patients that have already diagnosed. Such questionnaires cover several question fields with the focus of significant symptoms and binary classification as doctor diagnose result. The combination of question/answer pattern provides strong evidence for a particular disease.

In previous publishings of our study, we could manifest a fusion classifier diagnostic system [7] that unites and channelizes the qualitative benefits of single classifiers, i.e. support vector machines (SVM) [2], linear discriminant analysis (LDA) [12], logistic regression (LR) [12] and random forests (RF) [8]. The core of this system relies on methods that focus on assisting medical diagnostics [9], especially the previously migrated fusion method takes advantages of each classifier through evaluating their compatibility and accuracy of each derived diagnosis. Experimental approaches revealed that the fusion recognition, as a combination of various classifiers, presents a definite performance improvement rather than any single classifier outcomes [13]. For the maximal accurate diagnoses achievement, an continuous system improve is of central importance and will be the major focus within this paper.

Each classifier provides probabilities for the considered diseases. Some of the classifiers predict well-calibrated probabilities because they do not have biases. Some maximum margin algorithms such as SVM tend to push predicted probabilities away from 0 and 1, which brings characteristic sigmoid shaped distortion [3]. Other models such as Naive Bayes deliver unrealistic independence assumptions and tend to push probabilities closer to 0 and 1. For applications with different classifiers properties, it is essential to correct this bias gap and consider well-calibrated probabilities in case of multiple classifiers scoring scales [11].

Isotonic regression is a powerful calibration method that can help correcting monotonic distortions and rescale classifier probabilities into the same range, as the example shows in Figure 1. The calibration module allows us to obtain

better comparable probabilities of a given model, therefore increasing support to probability prediction. This makes it easier for the decision support system to rank examples in order of class-membership likelihood and find a more accurate probability of new patient that belongs to a particular class, thus reach a higher accuracy of decision making [15]. Another reason to do a calibration on disease classification scores is the positive samples are much less than negative ones. It keeps data in balance by introducing bias and make valuable data violent [11].

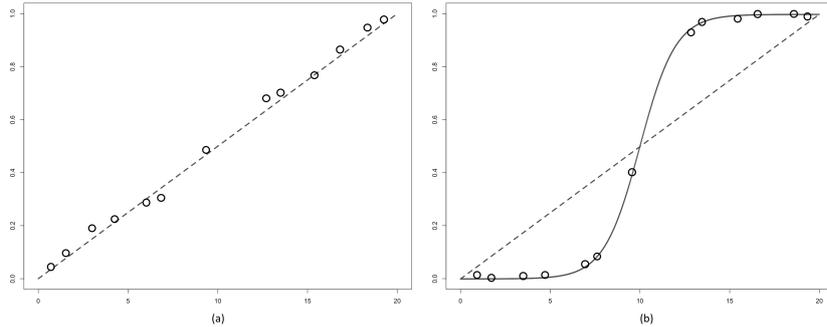


Fig. 1: Original Probability (a) vs. Calibrated Probability(b)

2 Related Work

This work relies on previous achievements and related studies with particular focus on multiple classifier systems, and best practice for the system-based decision improvements using probability scaling. The following concepts were methodologically adapted for the targeted medical diagnostic environment.

Zadrozny et al. [15] introduced the Pool Adjacent Violators Algorithm (PAVA) that enables the calibration of multi-class probability estimates in applying a simple ranked mapping methodology on classifier scores. The probability is an essential factor which represents the confidence level on the predicted outcome.

Niculescu-Mizil et al. [11] examine the probabilities predicted by ten supervised learning algorithms and the effectiveness of isotonic regression for calibrating the predictions made by different learning methods. It is shown that after calibration, most classifier models can predict better probability estimates.

Kortum et al. [7] investigated the benefits of using the classifier fusion method for diagnosing rare disease in practice. Experimental results show that this strategy dramatically improves the accuracy of the system compared to any single classifier. An auxiliary tool for physicians was realized that could derive computer-aided diagnoses in comparing a new patients symptoms with classified records from a shared patient database.

Chen et al. [3] further investigate the characteristic problem of classifier scores in medical diagnostic approaches. The authors proved that classifier scores on an

arbitrary scale could be converted to the probability scale for a target population prevalence value without affecting discrimination performance. This result takes an essential role in this paper since the dedicated probability problem addresses an interpretability gap, which solves a potential risk in classification.

3 Diagnostic Process for New Patient Data

In this chapter, we describe the classification improvement process for our computer aided diagnosis system that was originally revealed in our previous study stages [6,7]. Figure. 2 provides a compact structure overview about the decisioning architecture when classifying a new patient’s symptoms while considering probability calibration methods, e.g. the isotonic regression.

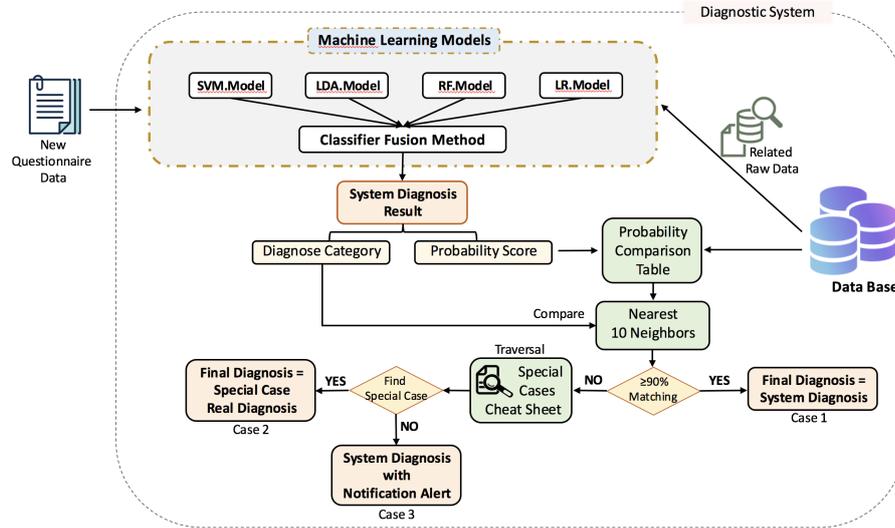


Fig. 2: System Diagnosis Process for a New Patient

When a new patient questionnaire data is available for diagnosis, the system’s first step is to extract the questionnaire related raw data from the database as the foundation of classification models. The classifier fusion method involves four algorithms (Support Vector Machine; Linear Discriminate Analysis; Random Forest and Logistic Regression) as a combination of multiple classifiers with robust prediction capabilities on medical data records in supervised modeling approaches [7]. After training each classification model by raw data, the system proceeds in predicting the probability of new coming data from each classifier. The fusion method is applied to obtain the diagnostic result due to the highest corresponding probability score, denoted as P .

Then the acquired P and its associated classification result will be compared within the range of its closest K -neighbors in the Probability Comparison Table (described in Section 3.3), K depends on the dataset size and the degree of dispersion of data within this interval. The comparison derives how many real doctors' decisions in a similar case are matching the obtained system diagnostic result. If the prediction presents a strong diagnosis probability ($\geq 90\%$ matched), the system decision will be selected as the final result. In case of a weak prediction value ($< 90\%$), the new patient symptoms data will be sequentially compared with the special case cheat sheet to check if there are any similar answer patterns in the database records with a firm diagnosis. A decisive matching (Case 2) will lead to the final diagnose according to the real doctors' diagnosis of the matched cases. If it does not find any matches in the cheat sheet (Case 3), a prompt message will be given along with the system diagnosis, indicating which variable affects the system to make the correct diagnosis because of insufficient evidence, thus helps doctors to examine the related symptoms tentatively. Compared with the initial diagnosis system, the subsequent process presents an additional improvement in filtering different cases between strong statistical predictions and poor predictions. This leads the system for more reliable results.

3.1 Classifier Fusion Method

The concept of applying multiple classifiers relies on the idea that the fusion method takes advantages of each classifier's outcome through evaluating their compatibility and accuracy, to derive the optimal computer-assisted diagnosis. Our previous studies proved that classifiers ensembles for symptoms interpretation, especially the combination of different supervised learning algorithms as shown in Figure 2, present strong performance improvements compared with any single classifier [7, 11].

The diagnostic process assigns a new dataset of a patient's disease symptoms to the trained classification pattern. The classifier fusion method applied in this study derives the average score from four supervised classifier scores $P(average)$. Each represents the likelihood for a particular type of disease d within the continuous range from 0 to 1. Further, each system-aided diagnosis d obtains the average probability value $P(d)$ as a representative indicator of a particular class. The diagnostic class from the highest $P(average)$ becomes selected as the diagnosis outcome. By evaluating the compatibility and accuracy of individual classifiers, the fusion method takes advantages of the single classifiers [9]. In case that one classifier is more particular about a specific diagnosis, it will occupy a more significant proportion within the fusion method. This allows a medical practitioner to derive more accurate diagnosis in comparing a new patients symptoms with the exact diagnostic records from the related dataset.

3.2 Isotonic Regression for Calibrating Classifier Scores

The classifier fusion model performs its system-based diagnoses through statistical classifiers; thus the outcome can generally be found within an arbitrary

scale. Toward that, the medical diagnostic meaning of classified decision scores is not always easy to interpret [15]. Modeling that involves the concept of calibration provides advantages when transferring derived classifier scores to a more meaningful probability space. The diagnosis of a disease is more reliable when the doctor can directly interpret the prediction score.

Therefore, the unnormalized scores produced from classifiers need to be calibrated to score-conditional probabilities, to supply reliable performance measures for generalization in the medical field [3]. For example, a patient’s classified diagnosis score of 0.80 presents a high likely indication for a particular disease – according to the average of all previously experienced patients. However, a patient with a classifier score of 0.45, on the other hand, does not clearly distinguish between types of diseases without straight comparison to other classified patient records. This leads to the concern about the discrimination and calibration of probabilities within multi-classifier systems [5].

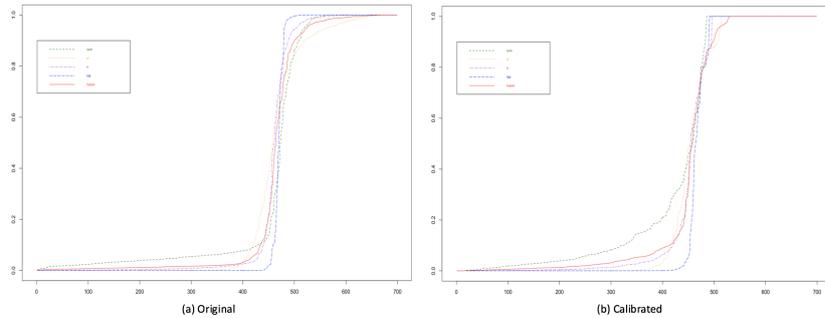


Fig. 3: Original Classifier Scores vs. Calibrated Scores - Dataset: BC

Figure 3 shows the difference between original classifier scores (a) and calibrated classifier scores (b), by using the Wisconsin breast cancer dataset from the UCI repository [4]. The calibrated classifier scores (b) tend to interpret the probability more closely as binary scoring without manipulating the models’ posterior probabilities. This can be improved with the help of an optimal threshold selection for binning the scores. The initial maximum margin of methods such as SVM push the posterior probability away from 0 and 1 while techniques such as LDA tend to push the probability towards 0 and 1. Methods like RF and SVM perform much better after applying isotonic regression for calibration, which helps to transform classification output into a probability distribution. Isotonic regression is a non-parametric technique that does not make any assumptions such as linearity among variables and constant error variance [14].

3.3 Integrated Probability Comparison Table

The Probability Comparison Table can be obtained by applying 10-Fold-Cross-Validation to traverse the entire raw data. As shown in Figure 4, the operation process is divided into three steps: Classifier Model Training; Calibration Model Training and Fuse Prediction Scores.

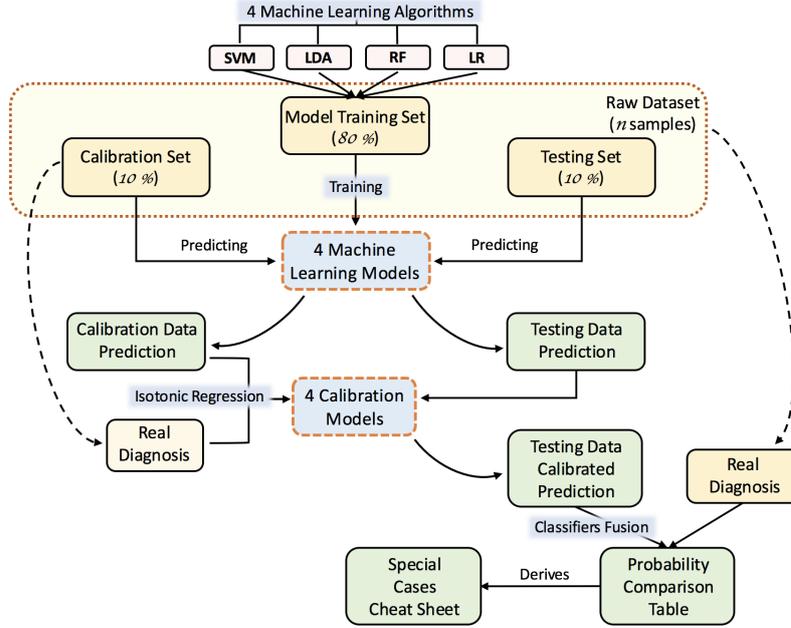


Fig. 4: Fusion Classifier Calibration for Minimized Classification Errors

In each cross-validation iteration, the applied dataset with one columns' binary classification will be partitioned into three parts: Model Training Set (80%); Calibration Set(10%) and Testing Set (10%). As the first step, 80% of the data records are used for training the classification models. The calibration model needs two elements to train: a prediction score and a corresponding real diagnosis. Thence before training the calibration model, the Calibration Set without diagnosis column will go through model prediction process to obtain the corresponding predictive value. By using isotonic regression along with the previously detached diagnosis column, the regular prediction values are used to train the calibration models. Probability calibration methods take commonly place for scaling each classifier scores' range into a population prevalences value [15]. This step is analogous to the applied testing, the alternative testing data prediction, except that its probability outcome has improved statistical characteristics due to the applied isotonic regression.

As a testing set that separates the doctors' original diagnosis column beside, it needs to be applied to four trained classification models and four calibration models to get calibrated prediction probabilities finally. After applying the fusion method [6, 7], the Probability Comparison Table will be generated, for later comparison and examination with the final system diagnosis. The table contains each questionnaire's original classifier scores, corresponding calibrated probability estimation, and real doctors' diagnoses of similar patient cases.

4 Evaluation of Disease Diagnostic System

We applied leave-one-out cross-validation (LOOCV) [1] to evaluate how well the improved model in Figure 2 performs on new data, by determining the diagnostic accuracy for patients correctly interpreted disease symptoms. Both systems, the initial and the improved become compared by their qualitative overall accuracy. When it comes to supervised learning algorithms, LOOCV presents a common strategy to quantify the system's decisional accuracy. We evaluated the system with three different datasets:

- Breast cancer dataset (*BC*) with 699 records [4]
- Primary immunodeficiency disorders dataset (*PID*) with 126 records [10]
- Rare disease dataset (*RD*) with 1021 records [6]

All the data records consist of the same characteristic: multiple attribute columns that describe the patients' symptoms condition and a binary diagnosis column given by a real doctor. The data we use in this paper are from patients that have been individually examined, tested and diagnosed by medical doctors.

The procedure of LOOCV assesses our model quality iteratively, by removing one sample from the dataset and sets it as testing data. The remaining samples are then used for training the model and predict the diagnosis category of the sample data that have been left out. After all data records have been sequentially extracted and diagnosed by the system, the model's overall accuracy can be calculated by the percentage of correct diagnosis compared with real doctors' decisions. Table 1 shows the system accuracy for the three different datasets. The second column lists the systems overall accuracy. The last column, which is the original diagnostic system accuracy for a classifier fusion without the use of isotonic regression; it can be seen that the strategy shown in Figure 2 can significantly enhance the overall accuracy of our decision support system.

The improvement for diagnosing breast cancer disease accurately only increases by 1%, foremost because the initial system diagnostic through the fusion classifier could already make a reliable diagnosis accuracy 96.1%. The other two datasets achieve a better improve with 3.6% for the primary immunodeficiency disorders recognition and 3.1% for rare disease diagnosis. Experimental proofs that all three diseases overall accuracy have been improved, whereas every improvement of system accuracy can provide influential decision support for the doctor's final diagnosis, thus helps doctors to apply proper treatments.

Data Set	Improved System Accuracy	Case 1 Accuracy	Case 2 Accuracy	Case 3 Accuracy	Original System Accuracy
BC	97.1%	98.6%	100%	92.7%	96.1%
PID	88.9%	95.2%	99.8%	71.7%	85.3%
RD	86.4%	93.5%	99.1%	66.6%	83.3%

Table 1: Disease Diagnostic System Overall Accuracy for three Datasets

In addition to the overall accuracy of the central diagnostic system, another essential system verification information is represented by the different accuracy of the three cases. Case 1 ensures the high prediction accuracy by comparing the system diagnosis results with the doctors’ decision in several similar cases, which required to achieve at least 90% matches. In the second case, the diagnosis accuracy could reach 99%. This reveals that the diagnostic result of some patients that have a relatively weak matching degree in case 1, could be adjusted by matching their data records with particular response patterns in which the system cannot make an accurate decision. Due to the similar answer patterns as in the reference records, the patient’s diagnosis result will follow the doctors’ suggestions in the reference data, which are substantiated correct in most case. Mainly, if there is a one-to-one correspondence between a new patient’s answer pattern and the special case in the cheat sheet, the system will ensure that the final diagnostic results are consistent with the doctor’s decision stored in the database, regardless of the system diagnosis. Case 3 mainly collects data records that the system cannot diagnose accurately. It is possible that all classifiers point to one diagnosis category, but the reality is poles apart. In this case, the best thing we can do is to mark the variables that influence the system diagnosis and send some notification to provide constructive suggestions to doctors for examining related symptoms.

5 Conclusion and Future Perspectives

Motivated by continuously improving the disease diagnostic system, we embedded the calibration procedure into the classifier fusion method. The conventional classification model for interpreting patient symptoms can only statistically refer to a particular type of disease. The newly integrated calibration adaption allow us to explain single symptom more accurately, due to the isotonic regression and derive diagnoses within three degrees of qualified outcomes (firm statistical determination, answer pattern matching and alert on unsure system decision).

We validated our updated diagnostic system by using three datasets covering 1846 patient records of multiple symptoms description of breast cancer, primary immunodeficiency disorders, rare diseases datasets with real doctors’ diagnoses. We compared the overall accuracy of the original fusion classifier model and the calibrated classification and proved the new strategy can improve system

performance. The system can help doctors to derive the correct diagnosis. More importantly, it can filter out the poor diagnostic outcome. We believe that the fusion classifier calibration and its case differentiation could help scientists, researchers and medical doctors to improve their patients' symptoms interpretation also in other data analysis areas. It presents an improvement measure for adequate diagnosis, especially for patients with symptom tendencies where the decidability between two types of diseases is not unmistakably apparent.

References

1. Arlot, S., Celisse, A.: A survey of cross-validation procedures for model selection. *Developments in Applied Statistics* 4 (2010) 4079, 159–177 (2003)
2. Auria, L., Moro, R.A.: Support vector machines (SVM) as a technique for solvency analysis (2008)
3. Chen, W., Sahiner, B., Samuelson, F., Pezeshk, A., Petrick, N.: Calibration of medical diagnostic classifier scores to the probability of disease. *Statistical methods in medical research* 27(5), 1394–1409 (2018)
4. Dr. William H. Wolberg, W. Nick Street, O.L.M.: UCI machine learning repository: Breast cancer wisconsin (1995), <http://archive.ics.uci.edu/ml/datasets>
5. H., S.C., L., G.J.: *Multivariate Classification Rules: Calibration and Discrimination*. American Cancer Society (2005)
6. Kortum, X., Grigull, L., Lechner, W., Klawonn, F.: A dynamic adaptive questionnaire for improved disease diagnostics. In: *Advances in Intelligent Data Analysis XVI*. pp. 162–172. Springer International Publishing (2017)
7. Kortum, X., Grigull, L., Muecke, U., Lechner, W., Klawonn, F.: Diagnosis support for orphan diseases. In: *International Conference on Intelligent Data Engineering and Automated Learning*. pp. 379–385. Springer International Publishing (2016)
8. Liaw, A., Wiener, M.: Classification and regression by random forest. *R news* 2(3), 18–22 (2002)
9. Ma, L., Liu, X., Song, L., Zhou, C., Zhao, X., Zhao, Y.: A new classifier fusion method based on historical and on-line classification reliability for recognizing common ct imaging signs of lung diseases. *Computerized Medical Imaging and Graphics* 40, 39–48 (2015)
10. Mücke, U., Klemann, C., Baumann, U., Meyer-Bahlburg, A., Kortum, X., Klawonn, F., Lechner, W.M., Grigull, L.: Patients experience in pediatric primary immunodeficiency disorders: computerized classification of questionnaires. *Frontiers in immunology* 8, 384 (2017)
11. Niculescu-Mizil, A., Caruana, R.: Predicting good probabilities with supervised learning. In: *Proceedings of the 22nd international conference on Machine learning*. pp. 625–632. ACM (2005)
12. Pohar, M., Blas, M., Turk, S.: Comparison of logistic regression and linear discriminant analysis: a simulation study. *Metodoloski zvezki* 1(1), 143 (2004)
13. Sboner, A., Eccher, C., Blanzieri, E., Bauer, P., Cristofolini, M., Zumiani, G., Forti, S.: A multiple classifier system for early melanoma diagnosis. *Artificial intelligence in medicine* 27(1), 29–44 (2003)
14. Stout, Q.F.: Isotonic regression algorithms. Retrieved: Aug. 6th (2011)
15. Zadrozny, B., Elkan, C.: Transforming classifier scores into accurate multiclass probability estimates. In: *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 694–699. ACM (2002)