

Chapter 1.25

Evaluation of Data Mining Methods

Paolo Giudici

University of Pavia, Italy

INTRODUCTION

Several classes of computational and statistical methods for data mining are available. Each class can be parameterised so that models within the class differ in terms of such parameters (see, for instance, Giudici, 2003; Hastie et al., 2001; Han & Kamber, 2000; Hand et al., 2001; Witten & Frank, 1999): for example, the class of linear regression models, which differ in the number of explanatory variables; the class of Bayesian networks, which differ in the number of conditional dependencies (links in the graph); the class of tree models, which differ in the number of leaves; and the class multi-layer perceptrons, which differ in terms of the number of hidden strata and nodes. Once a class of models has been established the problem is to choose the “best” model from it.

BACKGROUND

A rigorous method to compare models is statistical hypothesis testing. With this in mind one can adopt a sequential procedure that allows a model

to be chosen through a sequence of pairwise test comparisons. However, we point out that these procedures are generally not applicable in particular to computational data mining models, which do not necessarily have an underlying probabilistic model and, therefore, do not allow the application of statistical hypotheses testing theory. Furthermore, it often happens that for a data problem it is possible to use more than one type of model class, with different underlying probabilistic assumptions. For example, for a problem of predictive classification it is possible to use both logistic regression and tree models as well as neural networks.

We also point out that model specification and, therefore, model choice is determined by the type of variables used. These variables can be the result of transformations or of the elimination of observations, following an exploratory analysis. We then need to compare models based on different sets of variables present at the start. For example, how do we compare a linear model with the original explanatory variables with one with a set of transformed explanatory variables?

The previous considerations suggest the need

for a systematic study of the methods for comparison and evaluation of data mining models.

MAIN THRUST

Comparison criteria for data mining models can be classified schematically into: criteria based on statistical tests, based on scoring functions, computational criteria, Bayesian criteria and business criteria.

Criteria Based on Statistical Tests

The first are based on the theory of statistical hypothesis testing and, therefore, there is a lot of detailed literature related to this topic. See, for example, a text about statistical inference, such as Mood, Graybill, & Boes (1991) and Bickel & Doksum (1977). A statistical model can be specified by a discrete probability function or by a probability density function, $f(x)$. Such model is usually left unspecified, up to unknown quantities that have to be estimated on the basis of the data at hand. Typically, the observed sample it is not sufficient to reconstruct each detail of $f(x)$, but can indeed be used to approximate $f(x)$ with a certain accuracy. Often a density function is parametric so that it is defined by a vector of parameters $\Theta = (\theta_1, \dots, \theta_l)$, such that each value θ of Θ corresponds to a particular density function, $p_\theta(x)$. In order to measure the accuracy of a parametric model, one can resort to the notion of distance between a model f , which underlies the data, and an approximating model g (see, for instance, Zucchini, 2000). Notable examples of distance functions are, for categorical variables: the entropic distance, which describes the proportional reduction of the heterogeneity of the dependent variable; the chi-squared distance, based on the distance from the case of independence; and the 0-1 distance, which leads to misclassification rates. For quantitative variables, the typical choice is the Euclidean distance, representing the distance between two

vectors in a Cartesian space. Another possible choice is the uniform distance, applied when nonparametric models are being used.

Any of the previous distances can be employed to define the notion of discrepancy of an statistical model. The discrepancy of a model, g , can be obtained comparing the unknown probabilistic model, f , and the best parametric statistical model. Since f is unknown, closeness can be measured with respect to a sample estimate of the unknown density f . A common choice of discrepancy function is the Kullback-Leibler divergence, which can be applied to any type of observations. In such context, the best model can be interpreted as that with a minimal loss of information from the true unknown distribution.

It can be shown that the statistical tests used for model comparison are generally based on estimators of the total Kullback-Leibler discrepancy; the most used is the log-likelihood score. Statistical hypothesis testing is based on subsequent pairwise comparisons of log-likelihood scores of alternative models. Hypothesis testing allows one to derive a threshold below which the difference between two models is not significant and, therefore, the simpler models can be chosen.

Therefore, with statistical tests it is possible make an accurate choice among the models. The defect of this procedure is that it allows only a partial ordering of models, requiring a comparison between model pairs and, therefore, with a large number of alternatives it is necessary to make heuristic choices regarding the comparison strategy (such as choosing among the forward, backward and stepwise criteria, whose results may diverge). Furthermore, a probabilistic model must be assumed to hold, and this may not always be possible.

Criteria Based on Scoring Functions

A less structured approach has been developed in the field of information theory, giving rise to criteria based on score functions. These criteria

5 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/evaluation-data-mining-methods/7651

Related Content

Ethical Dilemmas in Data Mining and Warehousing

Joseph A. Cazier and Ryan C. LaBrie (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 2841-2849).

www.irma-international.org/chapter/ethical-dilemmas-data-mining-warehousing/7805

Conceptual Modeling Solutions for the Data Warehouse

Stefano Rizzi (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 208-227).

www.irma-international.org/chapter/conceptual-modeling-solutions-data-warehouse/7642

Agent-Based Mining of User Profiles for E-Services

Pasquale De Meo, Giovanni Quattrone, Giorgio Terracina and Domenico Ursino (2005). *Encyclopedia of Data Warehousing and Mining* (pp. 23-27).

www.irma-international.org/chapter/agent-based-mining-user-profiles/10559

Algorithms for Data Mining

Tadao Takaoka, Nigel K.L. Pope and Kevin E. Voges (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications* (pp. 1301-1319).

www.irma-international.org/chapter/algorithms-data-mining/7700

Hierarchical Document Clustering

Benjamin C.M. Fung, Ke Wang and Martin Ester (2005). *Encyclopedia of Data Warehousing and Mining* (pp. 555-559).

www.irma-international.org/chapter/hierarchical-document-clustering/10659