

## 36. Data Mining Methods and Applications

In this chapter, we provide a review of the knowledge discovery process, including data handling, data mining methods and software, and current research activities. The introduction defines and provides a general background to data mining knowledge discovery in databases. In particular, the potential for data mining to improve manufacturing processes in industry is discussed. This is followed by an outline of the entire process of knowledge discovery in databases in the second part of the chapter.

The third part presents data handling issues, including databases and preparation of the data for analysis. Although these issues are generally considered uninteresting to modelers, the largest portion of the knowledge discovery process is spent handling data. It is also of great importance since the resulting models can only be as good as the data on which they are based.

The fourth part is the core of the chapter and describes popular data mining methods, separated as supervised versus unsupervised learning. In supervised learning, the training data set includes observed output values ("correct answers") for the given set of inputs. If the outputs are continuous/quantitative, then we have a regression problem. If the outputs are categorical/qualitative, then we have a classification problem. Supervised learning methods are described in the context of both regression and classification (as appropriate), beginning with the simplest case of linear models, then presenting more complex modeling with trees, neural networks, and support vector machines, and concluding with some methods, such as nearest neighbor, that are only for classification. In unsupervised learning, the training data set does not contain output values. Unsupervised learning methods are described under two categories: association rules and clustering. Association rules are appropriate for business applications where precise numerical

36.1	<b>The KDD Process</b> .....	653
36.2	<b>Handling Data</b> .....	654
36.2.1	Databases and Data Warehousing	654
36.2.2	Data Preparation .....	654
36.3	<b>Data Mining (DM) Models and Algorithms</b> .....	655
36.3.1	Supervised Learning .....	655
36.3.2	Unsupervised Learning .....	661
36.3.3	Software .....	663
36.4	<b>DM Research and Applications</b> .....	664
36.4.1	Activity Monitoring .....	664
36.4.2	Mahalanobis–Taguchi System .....	665
36.4.3	Manufacturing Process Modeling .....	665
36.5	<b>Concluding Remarks</b> .....	667
	<b>References</b> .....	667

data may not be available while clustering methods are more technically similar to the supervised learning methods presented in this chapter. Finally, this section closes with a review of various software options.

The fifth part presents current research projects, involving both industrial and business applications. In the first project, data is collected from monitoring systems, and the objective is to detect unusual activity that may require action. For example, credit card companies monitor customers' credit card usage to detect possible fraud. While methods from statistical process control were developed for similar purposes, the difference lies in the quantity of data. The second project describes data mining tools developed by Genichi Taguchi, who is well known for his industrial work on robust design. The third project tackles quality and productivity improvement in manufacturing industries. Although some detail is given, considerable research is still needed to develop a practical tool for today's complex manufacturing processes.

Finally, the last part provides a brief discussion on remaining problems and future trends.

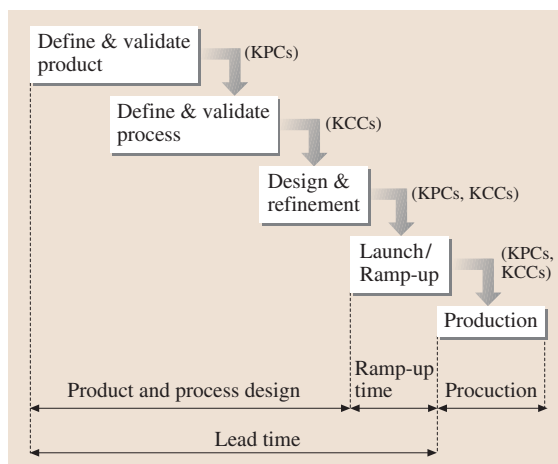
Data mining (DM) is the process of exploration and analysis, by automatic or semiautomatic means, of large quantities of data to discover meaningful patterns and rules [36.1]. Statistical DM is exploratory data analysis with little or no human interaction using computationally feasible techniques, i. e., the attempt to find *unknown interesting* structure [36.2]. Knowledge discovery in databases (KDD) is a multidisciplinary research field for nontrivial extraction of implicit, previously unknown, and potentially useful knowledge from data [36.3]. Although some treat DM and KDD equivalently, they can be distinguished as follows. The KDD process employs DM methods (algorithms) to extract knowledge according to the specifications of measures and thresholds, using a database along with any necessary preprocessing or transformations. DM is a step in the KDD process consisting of particular algorithms (methods) that, under some acceptable objective, produces particular patterns or knowledge over the data. The two primary fields that develop DM methods are statistics and computer science. Statisticians support DM by mathematical theory and statistical methods while computer scientists develop computational algorithms and relevant software [36.4]. Prerequisites for DM include: (1) Advanced computer technology (large CPU, parallel architecture, etc.) to allow fast access to large quantities of data and enable computationally intensive algorithms and statistical methods; (2) knowledge of the business or subject matter to formulate the important business questions and interpret the discovered knowledge.

With competition increasing, DM and KDD have become critical for companies to retain customers and ensure profitable growth. Although most companies are able to collect vast amounts of business data, they are often unable to leverage this data effectively to gain new knowledge and insights. DM is the process of applying sophisticated analytical and computational techniques to discover exploitable patterns in complex data. In many cases, the process of DM results in actionable knowledge and insights. Examples of DM applications include fraud detection, risk assessment, customer relationship management, cross selling, insurance, banking, retail, etc.

While many of these applications involve customer relationship management in the service industry, a potentially fruitful area is performance improvement and cost reduction through DM in industrial and manufacturing systems. For example, in the fast-growing and highly competitive electronics industry, total revenue worldwide in 2003 was estimated to be \$900 billion, and the growth rate is estimated at 8% per year

(www.selectron.com). However, economies of scale, purchasing power, and global competition are making the business such that one must either be a big player or serve a niche market. Today, extremely short life cycles and constantly declining prices are pressuring the electronics industry to manufacture their products with high quality, high yield, and low production cost.

To be successful, industry will require improvements at all phases of manufacturing. Figure 36.1 illustrates the three primary phases: design, ramp-up, and production. In the production phase, maintenance of a high performance level via improved system diagnosis is needed. In the ramp-up phase, reduction in new product development time is sought by achieving the required performance as quickly as possible. Market demands have been forcing reduced development time for new product and production system design. For example, in the computer industry, a product's life cycle has been shortened to 2–3 years recently, compared to a life cycle of 3–5 years a few years ago. As a result, there are a number of new concepts in the area of production systems, such as flexible and reconfigurable manufacturing systems. Thus, in the design phase, improved system performance integrated at both the ramp-up and production phases is desired. Some of the most critical factors and barriers in the competitive development of modern manufacturing systems lie in the largely uncharted area of predicting system performance during the design phase [36.5, 6]. Consequently, current systems necessitate that a large number of design/engineering changes be made after the system has been designed.



**Fig. 36.1** Manufacturing system development phases. KPCs = Key product characteristics. KCCs = Key control characteristics

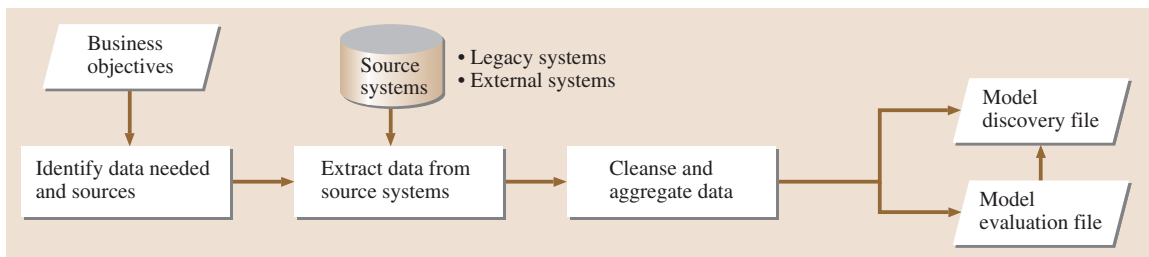
At all phases, system performance depends on many manufacturing process stages and hundreds or thousands of variables whose interactions are not well understood. For example, in the multi-stage printed circuit board (PCB) industry, the stages include process operations such as paste printing, chip placement, and wave soldering; and also include test operations such as optical inspection, vision inspection,

and functional test. Due to advancements in information technology, sophisticated software and hardware technologies are available to record and process huge amounts of daily data in these process and testing stages. This makes it possible to extract important and useful information to improve process and product performance through DM and quality improvement technologies.

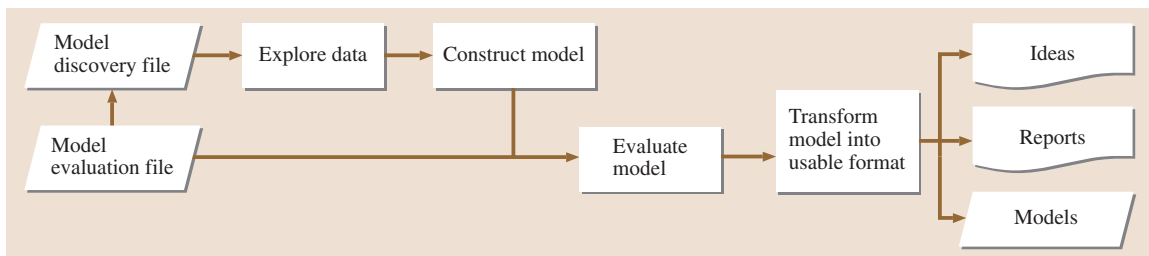
## 36.1 The KDD Process

The KDD process consists of four main steps:


1. Determination of business objectives,
2. Data preparation,
  - a) Create target data sets,
  - b) Data quality, cleaning, and preprocessing,
  - c) Data reduction and projection,
3. Data mining
  - a) Identify DM tasks,
  - b) Apply DM tools,
4. Consolidation and application,
  - a) Consolidate discovered knowledge,
  - b) Implement in business decisions.



**Fig. 36.2** Data preparation flow chart



**Fig. 36.3** Data mining flow chart



**Fig. 36.4** Consolidation and application flow chart

As an example of formulating business objectives, consider a telecommunications company. It is critically important to identify those customer traits that retain profitable customers and predict fraudulent behavior, credit risks and customer churn. This knowledge may be used to improve programs in target marketing, marketing channel management, micro-marketing, and cross selling. Finally, continually updating this

knowledge will enable the company to meet the challenges of new product development effectively in the future. Steps 2–4 are illustrated in figs. 36.2–36.4. Approximately 20–25% of effort is spent on determining business objectives, 50–60% of effort is spent on data preparation, 10–15% of is spent on DM, and about 10% is spent on consolidation/application.

## 36.2 Handling Data

The largest percentage effort of the KDD process is spent on processing and preparing the data. In this section, common forms of data storage and tools for accessing the data are described, and the important issues in data preparation are discussed.

### 36.2.1 Databases and Data Warehousing

A *relational database* system contains one or more objects called tables. The data or information for the database are stored in these tables. Tables are uniquely identified by their names and are comprised of columns and rows. Columns contain the column name, data type and any other attributes for the column. Rows contain the records or data for the columns. The structured query language (**SQL**) is the communication tool for relational database management systems. SQL statements are used to perform tasks such as updating data in a database, or retrieving data from a database. Some common relational database management systems that use SQL are: Oracle, Sybase, Microsoft SQL Server, Access, and Ingres. Standard SQL commands, such as *Select*, *Insert*, *Update*, *Delete*, *Create*, and *Drop*, can be used to accomplish almost everything that one needs to do with a database.

A *data warehouse* holds local databases assembled in a central facility. A *data cube* is a multidimensional array of data, where each dimension is a set of sets representing domain content, such as time or geography. The dimensions are scaled categorically, for example, region of country, state, quarter of year, week of quarter. The cells of the cube contain aggregated measures (usually counts) of variables. To explore the data cube, one can *drill down*, *drill up*, and *drill through*. Drill down involves splitting an aggregation into subsets, e.g., splitting region of country into states. Drill up involves consolidation, i.e., aggregating subsets along a dimension. Drill through involves subsets crossing multiple sets, e.g., the user might investigate statistics within

a state subset by time. Other databases and tools include object-oriented databases, transactional databases, time series and spatial databases, online analytical processing (**OLAP**), multidimensional OLAP (**MOLAP**), and relational OLAP using extended SQL (**ROLAP**). See Chapt. 2 of *Han and Kamber* [36.7] for more details.

### 36.2.2 Data Preparation

The purpose of this step in the KDD process is to identify data quality problems, sources of noise, data redundancy, missing data, and outliers. Data quality problems can involve inconsistency with external data sets, uneven quality (e.g., if a respondent fakes an answer), and biased opportunistically collected data. Possible sources of noise include faulty data collection instruments (e.g., sensors), transmission errors (e.g., intermittent errors from satellite or internet transmissions), data entry errors, technology limitations errors, misused naming conventions (e.g., using the same names for different meanings), and incorrect classification.

Redundant data exists when the same variables have different names in different databases, when a raw variable in one database is a derived variable in another, and when changes in a variable over time are not reflected in the database. These irrelevant variables impede the speed of the KDD process because dimension reduction is needed to eliminate them. Missing data may be irrelevant if we can extract useful knowledge without imputing the missing data. In addition, most statistical methods for handling missing data may fail for massive data sets, so new or modified methods still need to be developed. In detecting outliers, sophisticated methods like the Fisher information matrix or convex hull peeling are available, but are too complex for massive data sets. Although outliers may be easy to visualize in low dimensions, high-dimensional outliers may not show up in low-dimensional projections.

Currently, clustering and other statistical modeling are used.

The data preparation process involves three steps: data cleaning, database sampling, and database reduction and transformation. Data cleaning includes removal of duplicate variables, imputation of missing values, identification and correction of data inconsistencies, identification and updating of stale data, and creating a unique record (case) identification (ID). Via database sampling, the KDD process selects appropriate parts of the databases to be examined. For this to work, the data must satisfy certain conditions (e.g.,

no systematic biases). The sampling process can be expensive if the data have been stored in a database system such that it is difficult to sample the data the way you want and many operations need to be executed to obtain the targeted data. One must balance a trade-off between the costs of the sampling process and the mining process. Finally, database reduction is used for data cube aggregation, dimension reduction, elimination of irrelevant and redundant attributes, data compression, and encoding mechanisms via quantizations, wavelet transformation, principle components, etc.

### 36.3 Data Mining (DM) Models and Algorithms

The DM process is illustrated in Fig. 36.5. In this process, one will start by choosing an appropriate class of models. To fit the best model, one needs to split the sample data into two parts: the training data and the testing data. The training data will be used to fit the model and the testing data is used to refine and tune the fitted model. After the final model is obtained, it is recommended to use an independent data set to evaluate the goodness of the final model, such as comparing the prediction error to the accuracy requirement. (If independent data are not available, one can use the cross-validation method to compute prediction error.) If the accuracy requirement is not satisfied, then one must revisit earlier steps to reconsider other classes of models or collect additional data.

Before implementing any sophisticated DM methods, data description and visualization are used for initial exploration. Tools include descriptive statistical measures for central tendency/location, dispersion/spread, and distributional shape and symmetry; class characterizations and comparisons using analytical approaches, attribute relevance analysis, and class discrimination and comparisons; and data visualization using scatter-plot matrices, density plots, 3-D stereoscopic scatter-plots, and parallel coordinate plots. Following this initial step, DM methods take two forms: supervised versus unsupervised learning. Supervised learning is described as textlearning with a teacher, where the *teacher* provides data with correct answers. For example, if we want to classify online shoppers as buyers or non-buyers using an available set of variables, our data would include actual instances of buyers and non-buyers for training a DM method. Unsupervised learning is described as

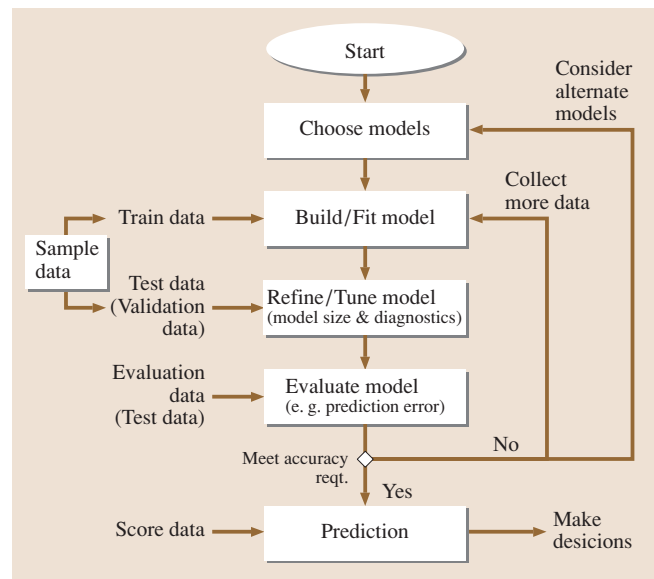


Fig. 36.5 Data mining process

*learning without a teacher*. In this case, correct answers are not available, and DM methods would search for patterns or clusters of similarity that could later be linked to some explanation.

#### 36.3.1 Supervised Learning

In supervised learning, we have a set of *input* variables (also known as predictors, independent variables,  $x$ ) that are measured or preset, and a set of *output* variables (also known as responses, dependent variables,  $y$ ) that are measured and assumed to be influenced by the in-

puts. If the outputs are continuous/quantitative, then we have a *regression* or *prediction* problem. If the outputs are categorical/qualitative, then we have a *classification* problem. First, a DM model/system is established based on the collected input and output data. Then, the established model is used to predict output values at new input values. The predicted values are denoted by  $\hat{y}$ .

The DM perspective of *learning with a teacher*, follows these steps:

- Student presents an answer ( $\hat{y}_i$  given  $\mathbf{x}_i$ );
- Teacher provides the correct answer  $y_i$  or an error  $\mathbf{e}_i$  for the student's answer;
- The result is characterized by some *loss function* or *lack-of-fit criterion*:  $\text{LOF}(\mathbf{y}, \hat{\mathbf{y}})$ ;
- The objective is to minimize the expected loss.

Supervised learning includes the common engineering task of function approximation, in which we assume that the output is related to the input via some function  $f(\mathbf{x}, \epsilon)$ , where  $\epsilon$  represents a random error, and seek to approximate  $f(\cdot)$ .

Below, we describe several supervised learning methods. All can be applied to both the regression and classification cases, except for those presented below under Other Classification Methods. We maintain the following notation. The  $j$ -th input variable is denoted by  $x_j$  (or random variable  $X_j$ ) and the corresponding boldface  $\mathbf{x}$  (or  $\mathbf{X}$ ) denotes the vector of  $p$  input variables  $(x_1, x_2, \dots, x_p)^T$ , where boldface  $\mathbf{x}_i$  denotes the  $i$ -th sample point;  $N$  is the number of sample points, which corresponds to the number of observations of the response variable; the response variable is denoted by  $y$  (or random variable  $Y$ ), where  $y_i$  denotes the  $i$ -th response observation. For the regression case, the response  $y$  is quantitative, while for the classification case, the response values are indices for  $C$  classes ( $c = 1, \dots, C$ ). An excellent reference for these methods is *Hastie et al.* [36.8].

### Linear and Additive Methods

In the regression case, the basic linear method is simply the *multiple linear regression* model form

$$\mu(\mathbf{x}; \beta) = E[Y | \mathbf{X} = \mathbf{x}] = \beta_0 + \sum_{m=1}^M \beta_m b_m(\mathbf{x}),$$

where the model terms  $b_m(\mathbf{x})$  are pre-specified functions of the input variables, for example, a simple linear term  $b_m(\mathbf{x}) = x_j$  or a more complex interaction term  $b_m(\mathbf{x}) = x_j x_k^2$ . The key is that the model is *linear in the parameters*  $\beta$ . Textbooks that cover linear

regression are abundant (e.g., [36.9, 10]). In particular, *Neter et al.* [36.11] provides a good background on residual analysis, model diagnostics, and model selection using best subsets and stepwise methods. In model selection, insignificant model terms are eliminated; thus, the final model may be a subset of the original pre-specified model. An alternate approach is to use a shrinkage method that employs a penalty function to *shrink* estimated model parameters towards zero, essentially reducing the influence of less important terms. Two options are ridge regression [36.12], which uses the penalty form  $\sum \beta_m^2$ , and the lasso [36.13], which uses the penalty form  $\sum |\beta_m|$ .

In the classification case, linear methods generate *linear decision boundaries* to separate the  $C$  classes. Although a direct linear regression approach could be applied, it is known not to work well. A better method is *logistic regression* [36.14], which uses log-odds (or logit transformations) of the posterior probabilities  $\mu_c(\mathbf{x}) = P(Y = c | \mathbf{X} = \mathbf{x})$  for classes  $c = 1, \dots, C - 1$  in the form

$$\begin{aligned} \log \frac{\mu_c(\mathbf{x})}{\mu_C(\mathbf{x})} &= \log \frac{P(Y = c | \mathbf{X} = \mathbf{x})}{P(Y = C | \mathbf{X} = \mathbf{x})} \\ &= \beta_{c0} + \sum_{j=1}^p \beta_{cj} x_j, \end{aligned}$$

where the  $C$  posterior probabilities  $\mu_c(\mathbf{x})$  must sum to one. The decision boundary between class  $c < C$  and class  $C$  is defined by the hyperplane  $\{\mathbf{x} | \beta_{c0} + \sum \beta_{cj} x_j = 0\}$ , where the log-odds are zero. Similarly, the decision boundary between classes  $c \neq C$  and  $d \neq C$ , derived from the log-odds for classes  $c$  and  $d$ , is defined by  $\{\mathbf{x} | \beta_{c0} + \sum \beta_{cj} x_j = \beta_{d0} + \sum \beta_{dj} x_j\}$ . In the binary case ( $C = 2$ ), if we define  $\mu(\mathbf{x}) = P(Y = 1 | \mathbf{X} = \mathbf{x})$ , then  $1 - \mu(\mathbf{x}) = P(Y = 2 | \mathbf{X} = \mathbf{x})$ . The logit transformation is then defined as  $g(\mu) = \mu / (1 - \mu)$ .

Closely related to logistic regression is *linear discriminant analysis* [36.15], which utilizes exactly the same linear form for the log-odds ratio, and defines linear *discriminant functions*  $\delta_c(\mathbf{x})$ , such that  $\mathbf{x}$  is classified to class  $c$  if its maximum discriminant is  $\delta_c(\mathbf{x})$ . The difference between the two methods is how the parameters are estimated. Logistic regression maximizes the conditional likelihood involving the posterior probabilities  $P(Y = c | \mathbf{X})$  while linear discriminant analysis maximizes the full log-likelihood involving the unconditional probabilities  $P(Y = c, \mathbf{X})$ . More general forms of discriminant analysis are discussed below under Other Classification Methods.

Finally, it should be noted that the logistic regression model is one form of *generalized linear model*



(GLM) [36.16]. GLM forms convert what appear to be nonlinear models into linear models, using tools such as transformations (e.g., logit) or conditioning on nonlinear parameters. This then enables the modeler to use traditional linear modeling analysis techniques. However, real data often do not satisfy the restrictive conditions of these models.

Rather than using pre-specified model terms, as in a linear model, a *generalized additive model* (GAM) [36.17], provides a more flexible statistical method to enable modeling of nonlinear patterns in each input dimension. In the regression case, the basic GAM form is

$$\mu(\mathbf{x}) = \beta_0 + \sum_{j=1}^p f_j(x_j),$$

where the  $f_j(\cdot)$  are unspecified (smooth) univariate functions, one for each input variable. The additive restriction prohibits inclusion of any interaction terms. Each function is fitted using a *nonparametric regression* modeling method, such as running-line smoothers (e.g., *lowess*, [36.18]), smoothing splines or kernel smoothers [36.19–21]. In the classification case, an additive logistic regression model utilizes the logit transformation for classes  $c = 1, \dots, C-1$  as above

$$\begin{aligned} \log \frac{\mu_c(\mathbf{x})}{\mu_C(\mathbf{x})} &= \log \frac{P(Y = c | \mathbf{X} = \mathbf{x})}{P(Y = C | \mathbf{X} = \mathbf{x})} \\ &= \beta_0 + \sum_{j=1}^p f_j(x_j), \end{aligned}$$

where an additive model is used in place of the linear model. However, even with the flexibility of nonparametric regression, GAM may still be too restrictive. The following sections describe methods that have essentially no assumptions on the underlying model form.

### Trees and Related Methods

One DM decision tree model is *chi-square automatic interaction detection* (CHAID) [36.22, 23], which builds non-binary trees using a chi-square test for the classification case and an  $F$ -test for the regression case. The CHAID algorithm first creates categorical input variables out of any continuous inputs by dividing them into several categories with approximately the same number of observations. Next, input variable categories that are not statistically different are combined, while a Bonferroni  $p$ -value is calculated for those that are statistically different. The best split is determined by the smallest  $p$ -value. CHAID continues to select splits un-

til the smallest  $p$ -value is greater than a pre-specified significance level ( $\alpha$ ).

The popular *classification and regression trees* (CART) [36.24] utilize recursive partitioning (binary splits), which evolved from the work of *Morgan* and *Sonquist* [36.25] and *Fielding* [36.26] on analyzing survey data. CARTs have a forward stepwise procedure that adds model terms and backward procedure for pruning. The model terms partition the  $\mathbf{x}$ -space into disjoint hyper-rectangular regions via indicator functions:  $b^+(x; t) = 1\{x > t\}$ ,  $b^-(x; t) = 1\{x \leq t\}$ , where the *split-point*  $t$  defines the borders between regions. The resulting model terms are:

$$f_m(\mathbf{x}) = \prod_{l=1}^{L_m} b^{s_{l,m}}(x_{v(l,m)}; t_{l,m}), \quad (36.1)$$

where,  $L_m$  is the number of univariate indicator functions multiplied in the  $m$ -th model term,  $x_{v(l,m)}$  is the input variable corresponding to the  $l$ -th indicator function in the  $m$ -th model term,  $t_{l,m}$  is the split-point corresponding to  $x_{v(l,m)}$ , and  $s_{l,m}$  is  $+1$  or  $-1$  to indicate the direction of the partition. The CART model form is then

$$f(\mathbf{x}; \beta) = \beta_0 + \sum_{m=1}^M \beta_m f_m(\mathbf{x}). \quad (36.2)$$

The partitioning of the  $\mathbf{x}$ -space does not keep the parent model terms because they are redundant. For example, suppose the current set has the model term:

$$f_a(\mathbf{x}) = 1\{x_3 > 7\} \cdot 1\{x_4 \leq 10\},$$

and the forward stepwise algorithm chooses to add

$$\begin{aligned} f_b(\mathbf{x}) &= f_a(\mathbf{x}) \cdot 1\{x_5 > 13\} \\ &= 1\{x_3 > 7\} \cdot 1\{x_4 \leq 10\} \cdot 1\{x_5 > 13\}. \end{aligned}$$

Then the model term  $f_a(\mathbf{x})$  is dropped from the current set. Thus, the recursive partitioning algorithm follows a binary tree with the current set of model terms  $f_m(\mathbf{x})$  consisting of the  $M$  leaves of the tree, each of which corresponds to a different region  $R_m$ .

In the regression case, CART minimizes the squared error loss function,

$$\text{LOF}(\hat{f}) = \sum_{i=1}^N [y_i - \hat{f}(x_i)]^2,$$

and the approximation is a piecewise-constant function. In the classification case, each region  $R_m$  is classified into one of the  $C$  classes. Specifically, define the

proportion of class  $c$  observations in region  $R_m$  as

$$\hat{\delta}_{mc} = \frac{1}{N_m} \sum_{x_i \in R_m} 1\{y_i = c\},$$

where  $N_m$  is the number of observations in the region  $R_m$ . Then the observations in region  $R_m$  are classified into the class  $c$  corresponding to the maximum proportion  $\hat{\delta}_{mc}$ . The algorithm is exactly the same as for regression, but with a different loss function. Appropriate choices include minimizing the misclassification error (i. e., the number of misclassified observations), the Gini index,  $\sum_{c=1}^C \hat{\delta}_{mc}(1 - \hat{\delta}_{mc})$ , or the deviance  $\sum_{c=1}^C \hat{\delta}_{mc} \log(\hat{\delta}_{mc})$ .

The exhaustive search algorithms for CART simultaneously conduct variable selection ( $x$ ) and split-point selection ( $t$ ). To reduce computational effort, the *fast algorithm for classification trees* [36.27] separates the two tasks. At each existing model term (leaf of the tree),  $F$ -statistics are calculated for variable selection. Then linear discriminant analysis is used to identify the split-point. A version for logistic and Poisson regression was presented by Chaudhuri et al. [36.28].

The primary drawback of CART and FACT is a bias towards selecting higher-order interaction terms due to the property of keeping only the leaves of the tree. As a consequence, these tree methods do not provide robust approximations and can have poor prediction accuracy. Loh and Shih [36.29] address this issue for FACT with a variant of their classification algorithm called QUEST that clusters classes into superclasses before applying linear discriminant analysis. For CART, Friedman et al. [36.30] introduced to the statistics literature the concepts of *boosting* [36.31] and *bagging* [36.32] from the machine learning literature. The bagging approach generates many bootstrap samples, fits a tree to each, then uses their average prediction. In the framework of boosting, a model term, called a *base learner*, is a *small* tree with only  $L$  disjoint regions ( $L$  is selected by the user), call it  $B(\mathbf{x}, \mathbf{a})$ , where  $\mathbf{a}$  is the vector of tree coefficients. The boosting algorithm begins by fitting a small tree  $B(\mathbf{x}, \mathbf{a})$  to the data, and the first approximation,  $\hat{f}_1(\mathbf{x})$ , is then this first small tree. In the  $m$ -th iteration, residuals are calculated, then a small tree  $B(\mathbf{x}, \mathbf{a})$  is fitted to the residuals and combined with the latest approximation to create the  $m$ -th approximation:

$$\begin{aligned} \hat{f}_m(\mathbf{x}; \beta_0, \beta_1, \dots, \beta_m) &= \hat{f}_{m-1}(\mathbf{x}; \beta_0, \beta_1, \\ &\dots, \beta_{m-1}) + \beta_m B(\mathbf{x}, \mathbf{a}), \end{aligned}$$

where a line search is used to solve for  $\beta_m$ . The resulting boosted tree, called a *multiple additive regression tree*

(MART) [36.33], then consists of much lower-order interaction terms. Friedman [36.34] presents *stochastic gradient boosting*, with a variety of loss functions, in which a bootstrap-like bagging procedure is included in the boosting algorithm.

Finally, for the regression case only, *multivariate adaptive regression splines* (MARS) [36.35] evolved from CART as an alternative to its piecewise-constant approximation. Like CART, MARS utilizes a forward stepwise algorithm to select model terms followed by a backward procedure to prune the model. A univariate version (appropriate for additive relationships) was presented by Friedman and Silverman [36.36]. The MARS approximation bends to model curvature at *knot* locations, and one of the objectives of the forward stepwise algorithm is to select appropriate knots. An important difference from CART is that MARS maintains the parent model terms, which are no longer redundant, but are simply lower-order terms.

MARS model terms have the same form as (36.1), except the indicator functions are replaced with truncated linear functions,

$$b^+(x; t) = [+(x - t)]_+, \quad b^-(x; t) = [-(x - t)]_+,$$

where  $[q]_+ = \max(0, q)$  and  $t$  is an univariate knot. The search for new model terms can be restricted to interactions of a maximum order (e.g.,  $L_m \leq 2$  permits up through two-factor interactions). The resulting MARS approximation, following (36.2), is a continuous, piecewise-linear function. After selection of the model terms is completed, smoothness to achieve a certain degree of continuity may be applied.

Hastie et al. [36.8] demonstrate significant improvements in accuracy using MART over CART. For the regression case, comparisons between MART and MARS yield comparable results [36.34]. Thus, the primary decision between these two methods is whether a piecewise-constant approximation is satisfactory or if a continuous, smooth approximation would be preferred.

### Artificial Neural Networks

*Artificial neural network* (ANN) models have been very popular for modeling a variety of physical relationships (for a general introduction see Lippmann [36.37] or Haykin [36.38]; for statistical perspectives see White [36.39], Baron et al. [36.40], Ripley [36.23], or Cheng and Titterton [36.41]). The original motivation for ANNs comes from how *learning* strengthens connections along neurons in the brain. Commonly, an ANN model is represented by a diagram of nodes in vari-



ous layers with weighted connections between nodes in different layers (Fig. 36.6). At the input layer, the nodes are the input variables and at the output layer, the nodes are the response variable(s). In between, there is usually at least one *hidden* layer which induces flexibility into the modeling. *Activation functions* define transformations between layers (e.g., input to hidden). Connections between nodes can *feed back* to previous layers, but for supervised learning, the typical ANN is *feedforward* only with at least one hidden layer.

The general form of a feedforward ANN with one hidden layer and activation functions  $b_1(\cdot)$  (input to hidden) and  $b_2(\cdot)$  (hidden to output) is

$$f_c(\mathbf{x}; \mathbf{w}, \mathbf{v}, \boldsymbol{\theta}, \gamma_c) = b_2 \left[ \sum_{h=1}^H w_{hc} \cdot b_1 \left( \sum_{j=1}^p v_{jh} x_j + \theta_h \right) + \gamma_c \right], \quad (36.3)$$

where  $c = 1, \dots, C$  and  $C$  is the number of output variables,  $p$  is the number of input variables,  $H$  is the number of hidden nodes, the weights  $v_{jh}$  link input nodes  $j$  to hidden nodes  $h$  and  $w_{hc}$  link hidden nodes  $h$  to output nodes  $c$ , and  $\theta_h$  and  $\gamma_c$  are constant terms called bias nodes (like intercept terms). The number of coefficients to be estimated is  $(p+1)H + (H+1)C$ , which is often larger than  $N$ . The simplest activation function is a linear function  $b(z) = z$ , which reduces the ANN model in (36.3) with one response variable to a multiple linear regression equation. For more flexibility, the recommended activation functions between the input and hidden layer(s) are the S-shaped *sigmoidal* functions or

the bell-shaped *radial basis functions*. Commonly used sigmoidal functions are the logistic function

$$b(z) = \frac{1}{1 + e^{-z}}$$

and the hyperbolic tangent

$$b(z) = \tanh(z) = \frac{1 - e^{-2z}}{1 + e^{-2z}}.$$

The most common radial basis function is the Gaussian probability density function.

In the regression case, each node in the output layer represents a quantitative response variable. The output activation function may be either a linear, sigmoidal, or radial basis function. Using a logistic activation function from input to hidden and from hidden to output, the ANN model in (36.3) becomes

$$f_c(\mathbf{x}; \mathbf{w}, \mathbf{v}, \boldsymbol{\theta}, \gamma_c) = \left[ 1 + \exp \left( - \sum_{h=1}^H w_{hc} z_h + \gamma_c \right) \right]^{-1},$$

where for each hidden node  $h$

$$z_h = \left[ 1 + \exp \left( - \sum_{j=1}^p v_{jh} x_j + \theta_h \right) \right]^{-1}.$$

In the classification case with  $C$  classes, each class is represented by a different node in the output layer. The recommended output activation function is the *softmax* function. For output node  $c$ , this is defined as


$$b(z_1, \dots, z_C; c) = \frac{e^{z_c}}{\sum_{d=1}^C e^{z_d}}.$$

This produces output values between zero and one that sum to one and, consequently, permits the output values to be interpreted as posterior probabilities for a categorical response variable.

Mathematically, an ANN model is a nonlinear statistical model, and a nonlinear method must be used to estimate the coefficients (weights  $v_{jh}$  and  $w_{hc}$ , biases  $\theta_h$  and  $\gamma_c$ ) of the model. This estimation process is called network training. Typically, the objective is to minimize the squared error lack-of-fit criterion

$$\text{LOF}(\hat{f}) = \sum_{c=1}^C \sum_{i=1}^N [y_i - \hat{f}_c(\mathbf{x}_i)]^2.$$

The most common method for training is *backpropagation*, which is based on gradient descent. At each



**Fig. 36.6** Diagram of a typical artificial neural network for function approximation. The input nodes correspond to the input variables, and the output node(s) correspond to the output variable(s). The number of hidden nodes in the hidden layer must be specified by the user

iteration, each coefficient (say  $w$ ) is adjusted according to its contribution to the lack-of-fit

$$\Delta w = \alpha \frac{\partial(\text{LOF})}{\partial w},$$

where the user-specified  $\alpha$  controls the step size; see Rumelhart et al. [36.42] for more details. More efficient training procedures are a subject of current ANN research.

Another major issue is the network *architecture*, defined by the number of hidden nodes. If too many hidden nodes are permitted, the ANN model will overfit the data. Many model discrimination methods have been tested, but the most reliable is validation of the model on a testing data set separate from the training data set. Several ANN architectures are fitted to the training data set and then prediction error is measured on the testing data set. Although ANNs are generally flexible enough to model anything, they are computationally intensive, and a significant quantity of representative data is required to both fit and validate the model. From a statistical perspective, the primary drawback is the overly large set of coefficients, none of which provide any intuitive understanding for the underlying model structure. In addition, since the nonlinear model form is not motivated by the true model structure, too few training data points can result in ANN approximations with extraneous nonlinearity. However, given enough good data, ANNs can outperform other modeling methods.

### Support Vector Machines

Referring to the linear methods for classification described earlier, the decision boundary between two classes is a hyperplane of the form  $\{\mathbf{x} \mid \beta_0 + \sum \beta_j x_j = 0\}$ . The *support vectors* are the points that are most critical to determining the optimal decision boundary because they lie close to the points belonging to the other class. With *support vector machines (SVM)* [36.43, 44], the linear decision boundary is generalized to the more flexible form

$$f(\mathbf{x}; \boldsymbol{\beta}) = \beta_0 + \sum_{m=1}^M \beta_m g_m(\mathbf{x}), \quad (36.4)$$

where the  $g_m(\mathbf{x})$  are transformations of the input vector. The decision boundary is then defined by  $\{\mathbf{x} \mid f(\mathbf{x}; \boldsymbol{\beta}) = 0\}$ . To solve for the optimal decision boundary, it turns out that we do not need to specify the transformations  $g_m(\mathbf{x})$ , but instead require only the *kernel function* [36.21, 45]:

$$K(\mathbf{x}, \mathbf{x}') = \left[ g_1(\mathbf{x}), \dots, g_M(\mathbf{x}) \right], \left[ g_1(\mathbf{x}'), \dots, g_M(\mathbf{x}') \right].$$

Two popular kernel functions for SVM are polynomials of degree  $d$ ,  $K(\mathbf{x}, \mathbf{x}') = (1 + \langle \mathbf{x}, \mathbf{x}' \rangle)^d$ , and radial basis functions,  $K(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/c)$ .

Given  $K(\mathbf{x}, \mathbf{x}')$ , we maximize the following Lagrangian dual-objective function:

$$\begin{aligned} \max_{\alpha_1, \dots, \alpha_N} \quad & \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{i'=1}^N \alpha_i \alpha_{i'} y_i y_{i'} K(\mathbf{x}_i, \mathbf{x}_{i'}) \\ \text{s.t.} \quad & 0 \leq \alpha_i \leq \gamma, \text{ for } i = 1, \dots, N \text{ and} \\ & \sum_{i=1}^N \alpha_i y_i = 0, \end{aligned}$$

where  $\gamma$  is an SVM tuning parameter. The optimal solution allows us to rewrite  $f(\mathbf{x}; \boldsymbol{\beta})$  as

$$f(\mathbf{x}; \boldsymbol{\beta}) = \beta_0 + \sum_{i=1}^N \alpha_i y_i K(\mathbf{x}, \mathbf{x}_i),$$

where  $\beta_0$  and  $\alpha_1, \dots, \alpha_N$  are determined by solving  $f(\mathbf{x}; \boldsymbol{\beta}) = 0$ . The support vectors are those  $\mathbf{x}_i$  corresponding to nonzero  $\alpha_i$ . A smaller SVM tuning parameter  $\gamma$  leads to more support vectors and a smoother decision boundary. A testing data set may be used to determine the best value for  $\gamma$ .

The SVM extension to more than two classes solves multiple two-class problems. SVM for regression utilizes the model form in (36.4) and requires specification of a loss function appropriate for a quantitative response [36.8, 46]. Two possibilities are the  $\epsilon$ -insensitive function

$$V_\epsilon(e) = \begin{cases} 0 & \text{if } |e| < \epsilon, \\ |e| - \epsilon & \text{otherwise,} \end{cases}$$

which ignores errors smaller than  $\epsilon$ , and the *Huber* [36.47] function

$$V_H(e) = \begin{cases} e^2/2 & \text{if } |e| \leq 1.345, \\ 1.345|e| - e^2/2 & \text{otherwise,} \end{cases}$$

which is used in robust regression to reduce model sensitivity to outliers.

### Other Classification Methods

In this section, we briefly discuss some other concepts that are applicable to DM classification problems. The basic intuition behind a good classification method is derived from the *Bayes classifier*, which utilizes the posterior distribution  $P(Y = c | \mathbf{X} = \mathbf{x})$ . Specifically, if  $P(Y = c | \mathbf{X} = \mathbf{x})$  is the maximum over  $c = 1, \dots, C$ , then  $\mathbf{x}$  would be classified to class  $c$ .

*Nearest neighbor* (NN) [36.48] classifiers seek to estimate the Bayes classifier directly without specification of any model form. The  $k$ -NN classifier identifies the  $k$  closest points to  $\mathbf{x}$  (using Euclidean distance) as the *neighborhood* about  $\mathbf{x}$ , then estimates  $P(Y = c | \mathbf{X} = \mathbf{x})$  with the fraction of these  $k$  points that are of class  $c$ . As  $k$  increases, the decision boundaries become smoother; however, the *neighborhood* becomes less local (and less relevant) to  $\mathbf{x}$ . This problem of local representation is even worse in high dimensions, and modifications to the distance measure are needed to create a practical  $k$ -NN method for DM. For this purpose, *Hastie and Tibshirani* [36.49] proposed the discriminant adaptive NN distance measure to reshape the neighborhood adaptively at a given  $\mathbf{x}$  to capture the critical points to distinguish between the classes.

As mentioned earlier, linear discriminant analysis may be too restrictive in practice. *Flexible discriminant analysis* replaces the linear decision boundaries with more flexible regression models, such as GAM or MARS. *Mixture discriminant analysis* relaxes the assumption that that classes are more or less spherical in shape by allowing a class to be represented by multiple (spherical) clusters; see *Hastie et al.* [36.50] and *Ripley* [36.23] for more details.

*K-means clustering classification* applies the  $K$ -means clustering algorithm separately to the data for each of the  $C$  classes. Each class  $c$  will then be represented by  $K$  clusters of points. Consequently, non-spherical classes may be modeled. For a new input vector  $\mathbf{x}$ , determine the closest cluster, then assign  $\mathbf{x}$  to the class associated with that cluster.

*Genetic algorithms* [36.51, 52] use processes such as genetic combination, mutation, and natural selection in an optimization based on the concepts of natural evolution. One generation of models competes to pass on characteristics to the next generation of models, until the best model is found. Genetic algorithms are useful in guiding DM algorithms, such as neural networks and decision trees [36.53].

### 36.3.2 Unsupervised Learning

In unsupervised learning, correct answers are not available, so there is no clear measure of success. Success must be judged subjectively by the value of discovered knowledge or the effectiveness of the algorithm. The statistical perspective is to observe  $N$  vectors from the population distribution, then conduct direct inferences on the properties (e.g. relationship, grouping) of the population distribution. The number of variables or

attributes is often very high (much higher than that in supervised learning). In describing the methods, we denote the  $j$ -th variable by  $x_j$  (or random variable  $X_j$ ), and the corresponding boldface  $\mathbf{x}$  (or  $\mathbf{X}$ ) denotes the vector of  $p$  variables  $(x_1, x_2, \dots, x_p)^T$ , where boldface  $\mathbf{x}_i$  denotes the  $i$ -th sample point. These variables may be either quantitative or qualitative.

#### Association Rules

*Association rules* or *affinity groupings* seek to find associations between the values of the variables  $\mathbf{X}$  that provide knowledge about the population distribution. *Market basket analysis* is a well-known special case, for which the extracted knowledge may be used to link specific products. For example, consider all the items that may be purchased at a store. If the analysis identifies that items  $A$  and  $B$  are commonly purchased together, then sales promotions could exploit this to increase revenue.

In seeking these associations, a primary objective is to identify variable values that occur together with high probability. Let  $S_j$  be the set of values for  $X_j$ , and consider a subset  $s_j \subseteq S_j$ . Then we seek subsets  $s_1, \dots, s_p$  such that

$$P \left[ \bigcap_{j=1}^p (X_j \in s_j) \right] \quad (36.5)$$

is large. In market basket analysis, the variables  $\mathbf{X}$  are converted to a set of binary variables  $\mathbf{Z}$ , where each attainable value of each  $X_j$  corresponds to a variable  $Z_k$ . Thus, the number of  $Z_k$  variables is  $K = \sum |S_j|$ . If binary variable  $Z_k$  corresponds to  $X_j = v$ , then  $Z_k = 1$  when  $X_j = v$  and  $Z_k = 0$  otherwise. An *item set*  $\kappa$  is a realization of  $\mathbf{Z}$ . For example, if the  $Z_k$  represent the possible products that could be purchased from a store, then an item set would be the set of items purchased together by a customer. Note that the number of  $Z_k = 1$  in an item set is at most  $p$ . Equation (36.5) now becomes

$$P \left[ \bigcap_{k \in \kappa} (Z_k = 1) \right],$$

which is estimated by

$$T(\kappa) = \frac{\text{Number of observations for which item set } \kappa \text{ occurs}}{N}.$$

$T(\kappa)$  is called the *support* for the rule. We can select a lower bound  $t$  such that item sets with  $T(\kappa) > t$  would be considered to have large support.

Further knowledge may be extracted via the a priori algorithm [36.54] in the form of if-then statements. For an item set  $\kappa$ , the items with  $Z_k = 1$  would be partitioned into two disjoint item subsets such that  $A \cup B = \kappa$ . The association rule would be stated as “if  $A$ , then  $B$ ” and denoted by  $A \Rightarrow B$ , where  $A$  is called the *antecedent* and  $B$  is called the *consequent*. This rule’s support  $T(A \Rightarrow B)$  is the same as  $T(\kappa)$  calculated above, an estimate of the joint probability. The *confidence* or *predictability* of this rule is

$$C(A \Rightarrow B) = \frac{T(A \Rightarrow B)}{T(A)},$$

which is an estimate of the conditional probability  $P(B|A)$ . The *expected confidence* is the support of  $B$ ,  $T(B)$ , and an estimate for the unconditional probability  $P(B)$ . The *lift* is the ratio of the confidence over the expected confidence,

$$L(A \Rightarrow B) = \frac{C(A \Rightarrow B)}{T(B)},$$

which, if greater than 1, can be interpreted as the increased prevalence of  $B$  when associated with  $A$ . For example, if  $T(B) = 5\%$ , then  $B$  is estimated to occur unconditionally 5% of the time. If  $C(A \Rightarrow B) = 40\%$ , then given  $A$  occurs,  $B$  is estimated to occur 40% of the time. This results in a lift of 8, implying that  $B$  is 8 times more likely to occur if we know that  $A$  occurs.

### Cluster Analysis

The objective of cluster analysis is to partition the  $N$  observations of  $\mathbf{x}$  into groups or clusters such that the dissimilarities within each cluster are smaller than the dissimilarities between different clusters [36.55]. Typically the variables  $\mathbf{x}$  are all quantitative, and a distance measure (e.g., Euclidean) is used to measure dissimilarity. For categorical  $\mathbf{x}$  variables, a dissimilarity measure must be explicitly defined. Below, we describe some of the more common methods.

*K-means* [36.56] is the best-known clustering tool. It is appropriate when the variables  $\mathbf{x}$  are quantitative. Given a prespecified value  $K$ , the method partitions the  $N$  observations of  $\mathbf{x}$  into exactly  $K$  clusters by minimizing within-cluster dissimilarity. Squared Euclidean distance

$$d(\mathbf{x}_i, \mathbf{x}'_j) = \sum_{j=1}^p (x_{ij} - x'_{ij})^2$$

is used to measure dissimilarity. For a specific clustering assignment  $\mathbf{C} = (C_1, \dots, C_K)$ , the within-cluster

dissimilarity is measured by calculating  $d(\mathbf{x}_i, \mathbf{x}'_j)$  for all points  $\mathbf{x}_i, \mathbf{x}'_j$  within a cluster  $C_k$ , then summing over the  $K$  clusters. This is equivalent to calculating

$$W(\mathbf{C}) = \sum_{k=1}^K \sum_{i \in C_k} d(\mathbf{x}_i, \bar{\mathbf{x}}_k),$$

where the cluster mean  $\bar{\mathbf{x}}_k$  is the sample mean vector of the points in cluster  $C_k$ . Given a current set of cluster means, the  $K$ -means algorithm assigns each point to the closest cluster mean, calculates the new cluster means, and iterates until the cluster assignments do not change.

Unfortunately, because of its dependence on the squared Euclidean distance measure,  $K$ -means clustering is sensitive to outliers (i.e., is not robust). *K-medoids* [36.57] is a generalized version that utilizes an alternately defined cluster *center* in place of the cluster means and an alternate distance measure.

*Density-based clustering (DBSCAN)* [36.58] algorithms are less sensitive to outliers and can discover clusters of irregular ( $p$ -dimensional) shapes. DBSCAN is designed to discover clusters and noise in a spatial database. The advantage of DBSCAN over other clustering methods is its ability to represent specific structure in the analysis explicitly. DBSCAN has two key parameters: neighborhood size ( $\epsilon$ ) and minimum cluster size ( $n_{\min}$ ). The neighborhood of an object within a radius  $\epsilon$  is called the  $\epsilon$ -neighborhood of the object. If the  $\epsilon$ -neighborhood of an object contains at least  $n_{\min}$  objects, then the object is called a core object. To find a cluster, DBSCAN starts with an arbitrary object  $o$  in the database. If the object  $o$  is a core object w.r.t.  $\epsilon$  and  $n_{\min}$ , then a new cluster with  $o$  as the core object is created. DBSCAN continues to retrieve all density-reachable objects from the core object and add them to the cluster.

GDBSCAN [36.59] generalizes the two key parameters of the DBSCAN algorithm such that it can cluster point objects and spatially extended objects according to an arbitrarily selected combination of attributes. The neighborhood of an object is now defined by a binary predicate  $\eta$  on a data set that is reflexive and symmetric. If  $\eta$  is true, then the neighborhood of an object is called the  $\eta$ -neighborhood of an object. In other words, the  $\eta$ -neighborhood of an object is a set of objects,  $S$ , which meet the condition that  $\eta$  is true. Corresponding to  $n_{\min}$ , another predicate,  $w_{\min}$  of the set of objects,  $S$ , is defined such that it is true if and only if the weighted cardinality for the set,  $w\text{Card}(S)$ , is greater or equal to the minimum cardinality (MinCard), i.e.  $w\text{Card}(S) \geq \text{MinCard}$ .

Finally, ordering points to identify the clustering structure (OPTICS) [36.60] is a method of cluster analysis that produces an augmented ordering of the database representing its density-based clustering structure. This method by itself does not produce a clustering of a data set explicitly. The information produced by OPTICS includes representative points, arbitrarily shaped clusters and intrinsic clustering structure, which can then be used by a clustering algorithm when selecting clustering settings. This same information can also be used by a human expert to gain insight into the clustering structure of the data.

*Self-Organizing (Feature) Maps (SOMs)* [36.61] belong to the class of ANNs called *unsupervised learning networks*. SOMs can be organized as a single layer or as two layers of neuron nodes. In this arrangement, the input layer consists of  $p$  nodes corresponding to the real-valued input vector of dimension  $p$ . The input layer nodes are connected to a second layer of nodes  $U$ . By means of lateral connections, the nodes in  $U$  form a lattice structure of dimensionality  $M$ . Typically  $M$  is much smaller than  $p$ . By means of a learning algorithm, the network *discovers* the clusters within the data. It is possible to alter the discovered clusters by varying the learning parameters of the network. The SOM is especially suitable for data survey because it has appealing visualization properties. It creates a set of prototype vectors representing the data set and carries out a topology-preserving projection of the prototypes from the  $p$ -dimensional input space onto a low-dimensional (typically two-dimensional) grid. This ordered grid can be used as a convenient visualization surface for showing different features of the SOM (and thus of the data), for example, the cluster structure. While the axes of such a grid do not correspond to any measurement, the spatial relationships among the clusters do correspond to relationships in  $p$ -dimensional space. Another attractive feature of the SOM is its ability to discover arbitrarily shaped clusters organized in a nonlinear space.

### 36.3.3 Software

Several DM software packages are available at a wide range of prices, of which six of the most popular packages are:

- SAS Enterprise Miner ([www.sas.com/technologies/analytics/datamining/miner/](http://www.sas.com/technologies/analytics/datamining/miner/)),
- SPSS Clementine ([www.spss.com/clementine/](http://www.spss.com/clementine/)),
- XLMiner in Excel ([www.xlminer.net](http://www.xlminer.net)),

- Ghostminer ([www.fqspl.com.pl/ghostminer/](http://www.fqspl.com.pl/ghostminer/)),
- Quadstone ([www.quadstone.com/](http://www.quadstone.com/)),
- Insightful Miner ([www.splus.com/products/iminer/](http://www.splus.com/products/iminer/)).

Haughton et al. [36.62] present a review of the first five listed above. The SAS and SPSS packages have the most complete set of KDD/DM tools (data handling, DM modeling, and graphics), while Quadstone is the most limited. Insightful Miner was developed by  $S+$  [[www.splus.com](http://www.splus.com)], but does not require knowledge of the  $S+$  language, which is only recommended for users that are familiar with statistical modeling. For statisticians, the advantage is that Insightful Miner can be integrated with more sophisticated DM methods available with  $S+$ , such as flexible and mixture discriminant analysis. All six packages include trees and clustering, and all except Quadstone include ANN modeling. The SAS, SPSS, and XLMiner packages include discriminant analysis and association rules. Ghostminer is the only one that offers SVM tools.

SAS, SPSS, and Quadstone are the most expensive (over \$ 40 000) while XLMiner is a good deal for the price (under \$ 2 000). The disadvantage of XLMiner is that it cannot handle very large data sets. Each package has certain specializations, and potential users must carefully investigate these choices to find the package that best fits their KDD/DM needs. Below we describe some other software options for the DM modeling methods presented.

GLM or linear models are the simplest of DM tools and most statistical software can fit them, such as SAS, SPSS,  $S+$ , and Statistica [[www.statsoftinc.com/](http://www.statsoftinc.com/)]. However, it should be noted that Quadstone only offers a regression tool via scorecards, which is not the same as statistical linear models. GAM requires access to more sophisticated statistical software, such as  $S+$ .

Software for CART, MART, and MARS is available from Salford Systems [[www.salford-systems.com](http://www.salford-systems.com)]. SAS Enterprise Miner includes CHAID, CART, and the machine learning program C4.5 [[www.rulequest.com](http://www.rulequest.com)] [36.63], which uses classifiers to generate decision trees and if-then rules. SPSS Clementine and Insightful Miner also include CART, but Ghostminer and XLMiner utilize different variants of decision trees. QUEST [[www.stat.wisc.edu/loh/quest.html](http://www.stat.wisc.edu/loh/quest.html)] is available in SPSS's AnswerTree software and Statistica.

Although ANN software is widely available, the most complete package is Matlab's [[www.mathworks.com](http://www.mathworks.com)] Neural Network Toolbox. Information on SVM software is available at [[www.support-vector.net/software.html](http://www.support-vector.net/software.html)]. One good option is Matlab's SVM Toolbox.



## 36.4 DM Research and Applications

Many industrial and business applications require modeling and monitoring processes with real-time data of different types: real values, categorical, and even text. DM is an effective tool for extracting process knowledge and discovering data patterns to provide a control aid for these processes. Advanced DM research involves complex system modeling of heterogeneous objects, where adaptive algorithms are necessary to capture dynamic system behavior.

### 36.4.1 Activity Monitoring

One important DM application is the development of an effective data modeling and monitoring system for understanding customer profiles and detecting fraudulent behavior. This is generally referred to as *activity monitoring for interesting events requiring action* [36.64]. Other activity monitoring examples include credit card or insurance fraud detection, computer intrusion detection, some forms of fault detection, network performance monitoring, and news story monitoring.

Although activity monitoring has only recently received attention in the information industries, solutions to similar problems were developed long ago in the manufacturing industries, under the moniker *statistical process control (SPC)*. SPC techniques have been used routinely for online process control and monitoring to achieve process stability and to improve process capability through variation reduction. In general, all processes are subject to some natural variability regardless of their state. This natural variability is usually small and unavoidable and is referred to as *common cause variation*. At the same time, processes may be subject to other variability caused by improper machine adjustment, operator errors, or low-quality raw material. This variability is usually large, but avoidable, and is referred to as *special cause variation*. The basic objective of SPC is to detect the occurrence of special cause variation (or process shifts) quickly, so that the process can be investigated and corrective action may be taken before quality deteriorates and defective units are produced. The main ideas and methods of SPC were developed in the 1920s by Walter Shewhart of Bell Telephone Laboratories and have had tremendous success in manufacturing applications [36.65, 66]. *Montgomery* and *Woodall* [36.67] provide a comprehensive panel discussion on SPC, and multivariate methods are reviewed by *Hayter* and *Tsui* [36.68] and *Ma-son* et al. [36.69].

Although the principle of SPC can be applied to service industries, such as business process monitoring, fewer applications exist for two basic reasons that *Montgomery* [36.65] identified. First, the system that needs to be monitored and improved is obvious in manufacturing applications, while it is often difficult to define and observe in service industries. Second, even if the system can be clearly specified, most non-manufacturing operations do not have natural measurement systems that reflect the performance of the system. However, these obstacles no longer exist, due to the many natural and advanced measurement systems that have been developed. In the telecommunications industry, for example, advanced software and hardware technologies make it possible to record and process huge amounts of daily data in business transactions and service activities. These databases contain potentially useful information to the company that may not be discovered without knowledge extraction or DM tools.

While SPC ideas can be applied to business data, SPC methods are not directly applicable. Existing SPC theories are based on small or medium-sized samples, and the basic hypothesis testing approach is intended to detect only simple shifts in a process mean or variance. Recently, *Jiang* et al. [36.70] successfully generalized the SPC framework to model and track thousands of diversified customer behaviors in the telecommunication industry. The challenge is to develop an integrated strategy to monitor the performance of an entire multi-stage system and to develop effective and efficient techniques for detecting the systematic changes that require action.

A dynamic business process can be described by the dynamic linear models introduced by *West* [36.71],

$$\text{Observation equation: } X_t = A_t \theta_t + \Delta_t,$$

$$\text{System evolution equation: } \theta_t = B_t \theta_{t-1} + \Lambda_t,$$

$$\text{Initial information: } \pi(S_0),$$

where  $A_t$  and  $B_t$  represent observation and state transition matrices, respectively, and  $\Delta_t$  and  $\Lambda_t$  represent observation and system transition errors, respectively. Based on the dynamic system model, a model-based process monitoring and root-cause identification method can be developed. Monitoring and diagnosis includes fault pattern generation and feature extraction, isolation of the critical processes, and root-cause identification. *Jiang* et al. [36.70] utilize this for individual customer prediction and monitoring. In general, individual modeling is computationally intractable and

cluster models should be developed with mixture distributions [36.72].

One particularly competitive industry is telecommunications. Since divestiture and government deregulation, various telephone services, such as cellular, local and long distance, domestic and commercial, have become battle grounds for telecommunication service providers. Because of the data and information oriented nature of the industry, DM methods for knowledge extraction are critical. To remain competitive, it is important for companies to develop business planning systems that help managers make good decisions. In particular, these systems will allow sales and marketing people to establish successful customer loyalty programs for churn prevention and to develop fraud detection modules for reducing revenue loss through market segmentation and customer profiling.

A major task in this research is to develop and implement DM tools within the business planning system. The objectives are to provide guidance for targeting business growth, to forecast year-end usage volume and revenue growth, and to value risks associated with the business plan periodically. Telecommunication business services include voice and non-voice services, which can be further categorized to include domestic, local, international, products, toll-free calls, and calling cards. For usage forecasting, a minutes growth model is utilized to forecast domestic voice usage. For revenue forecasting, the average revenue per minute on a log scale is used as a performance measure and is forecasted by a double exponential smoothing growth function. A structural model is designed to decompose the business growth process into three major subprocesses: add, disconnect, and base. To improve explanatory power, the revenue unit is further divided into different customer groups. To compute confidence and prediction intervals, bootstrapping and simulation methods are used.

To understand the day effect and seasonal effect, the concept of bill-month equivalent business days (EBD) is defined and estimated. To estimate EBD, the factor characteristics of holidays (non-EBD) are identified and eliminated and the day effect is estimated. For seasonality, the US Bureau of the Census X-11 seasonal adjustment procedure is used.

### 36.4.2 Mahalanobis–Taguchi System

Genichi Taguchi is best known for his work on robust design and design of experiments. The Taguchi robust design methods have generated a considerable

amount of discussion and controversy and are widely used in manufacturing [36.73–77]. The general consensus among statisticians seems to be that, while many of Taguchi's overall ideas on experimental design are very important and influential, the techniques he proposed are not necessarily the most effective statistical methods. Nevertheless, Taguchi has made significant contributions in the area of quality control and quality engineering. For DM, Taguchi has recently popularized the *Mahalanobis–Taguchi System (MTS)*, a new set of tools for diagnosis, classification, and variable selection. The method is based on a Mahalanobis distance scale that is utilized to measure the level of abnormality in *abnormal* items as compared to a group of *normal* items. First, it must be demonstrated that a Mahalanobis distance measure based on all available variables is able to separate the abnormal from the normal items. Should this be successfully achieved, orthogonal arrays and signal-to-noise ratios are used to select an *optimal* combination of variables for calculating the Mahalanobis distances.

The MTS method has been claimed to be very powerful for solving a wide range of problems, including manufacturing inspection and sensing, medical diagnosis, face and voice recognition, weather forecasting, credit scoring, fire detection, earthquake forecasting, and university admissions. Two recent books have been published on the MTS method by *Taguchi et al.* [36.78] and *Taguchi and Jugulum* [36.79]. Many successful case studies in MTS have been reported in engineering and science applications in many large companies, such as Nissan Motor Co., Mitsubishi Space Software Co., Xerox, Delphi Automotive Systems, ITT Industries, Ford Motor Company, Fuji Photo Film Company, and others. While the method is getting a lot of attention in many industries, very little research [36.80] has been conducted to investigate how and when the method is appropriate.

### 36.4.3 Manufacturing Process Modeling

One area of DM research in manufacturing industries is quality and productivity improvement through DM and knowledge discovery. Manufacturing systems nowadays are often very complicated and involve many manufacturing process stages where hundreds or thousands of in-process measurements are taken to indicate or initiate process control of the system. For example, a modern semiconductor manufacturing process typically consists of over 300 steps, and in each step, multiple pieces of equipment are used

to process the wafer. Inappropriate understanding of interactions among in-process variables will create inefficiencies at all phases of manufacturing, leading to long product/process realization cycle times and long development times, and resulting in excessive system costs.

Current approaches to DM in electronics manufacturing include neural networks, decision trees, Bayesian models and rough set theory [36.81, 82]. Each of these approaches carries certain advantages and disadvantages. Decision trees, for instance, produce intelligible rules and hence are very appropriate for generating process control or design of experiments strategies. They are, however, generally prone to outlier and imperfect data influences. Neural networks, on the other hand, are robust against data abnormalities but do not produce readily intelligible knowledge. These methods also differ in their ability to handle high-dimensional data, to discover arbitrarily shaped clusters [36.58] and to provide a basis for intuitive visualization [36.83]. They can also be sensitive to training and model building parameters [36.60]. Finally, the existing approaches do not take into consideration the localization of process parameters. The patterns or clusters identified by existing approaches may include parameters from a diverse set of components in the system. Therefore, a combination of methods that complement each other to provide a complete set of desirable features is necessary.

It is crucial to understand process structure and yield components in manufacturing, so that problem localization can permit reduced production costs. For example, semiconductor manufacturing practice shows that over 70% of all fatal defects and close to 90% of yield excursions are caused by problems related to process equipment [36.84]. Systematic defects can be attributed to many categories that are generally associated with technologies and combinations of different process operations. To implement DM methods successfully for knowledge discovery, some future research for manufacturing process control must include yield modeling, defect modeling and variation propagation.

### Yield Modeling

In electronics manufacturing, the ANSI standards [36.85] and practice generally assume that the number of defects on an electronics product follows a Poisson distribution with mean  $\lambda$ . The Poisson random variable is an approximation of the sum of independent Bernoulli trials, but defects on different components may be correlated since process yield critically de-

pends on product groups, process steps, and types of defects [36.86]. Unlike traditional defect models, an appropriate logit model can be developed as follows. Let the number of defects of category  $X$  on an electronics product be

$$U_X = \sum Y_X$$

and

$$\begin{aligned} \text{logit}[E(Y_X)] = & \alpha_X^0 + \alpha_X^O \cdot O_X \\ & + \alpha_X^C \cdot C_X + \alpha_X^{OC} \cdot O_X \cdot C_X, \end{aligned}$$

where  $\text{logit}(z) = \log[z/(1-z)]$  is the link function for Bernoulli distributions, and  $Y_X$  is a Bernoulli random variable representing a defect from defect category  $X$ . The default logit of the failure probability is  $\alpha_X^0$ , and  $\alpha_X^O$  and  $\alpha_X^C$  are the main effects of operations ( $O_X$ ) and components ( $C_X$ ). Since the  $Y_X$ s are correlated, this model will provide more detailed information about defects.

### Multivariate Defect Modeling

Since different types of defects may be caused by the same operations, multivariate Poisson models are necessary to account for correlations among different types of defects. The *trivariate reduction method* suggests an additive Poisson model for the vector of Poisson counts  $U = (U_1, U_2, \dots, U_k)'$ ,

$$U = AV,$$

where  $A$  is a matrix of zeros and ones, and  $V = (v_1, v_2, \dots, v_p)'$  consists of independent Poisson variables  $v_i$ . The variance-covariance matrix takes the form  $\text{Var}(U) = A \Sigma A' = \Phi + \nu \nu'$ , where  $\Phi = \text{diag}(\mu_i)$  is a diagonal matrix with the mean of the individual series, and  $\nu$  is the common covariance term. Note that the  $v_i$  are essentially latent variables, and a factor analysis model can be developed for analyzing multivariate discrete Poisson variables such that

$$\log[E(U)] = \mu + L \cdot F,$$

where  $U$  is the vector of defects,  $L$  is the matrix of factor loadings, and  $F$  contains common factors representing effects of specific operations. By using factor analysis, it is possible to relate product defects to the associated packages and operations.

### Multistage Variation Propagation

Inspection tests in an assembly line usually have functional overlap, and defects from successive inspection

stations exhibit strong correlations. Modeling serially correlated defect counts is an important task for defect localization and yield prediction. Poisson regression models, such as the generalized event-count method [36.87] and its alternatives, can be utilized to

account for serial correlations of defects in different inspection stations. Factor analysis methods based on hidden Markov models [36.88] can also be constructed to investigate how variations are propagated through assembly lines.

## 36.5 Concluding Remarks

While DM and KDD methods are gaining recognition and have become very popular in many companies and enterprises, the success of these methods is still somewhat limited. Below, we discuss a few obstacles.

First, the success of DM depends on a close collaboration of subject-matter experts and data modelers. In practice, it is often easy to identify the right subject-matter expert, but difficult to find the qualified data modeler. While the data modeler must be knowledgeable and familiar with DM methods, it is more important to be able to formulate real problems such that the existing methods can be applied. In reality, traditional academic training mainly focuses on knowledge of modeling algorithms and lacks training in problem formulation and interpretation of results. Consequently, many modelers are very efficient in fitting models and algorithms to data, but have a hard time determining when and why they should use certain algorithms. Similarly, the ex-

isting commercial DM software systems include many sophisticated algorithms, but lack of guidance on which algorithms to use.

Second, implementation of DM is difficult to apply effectively across an industry. Although it is clear that extracting hidden knowledge and trends across an industry would be useful and beneficial to all companies in the industry, it is typically impossible to integrate the detailed data from competing companies due to confidentiality and proprietary issues. Currently, the industry practice is that each company will integrate their own detailed data with the more general, aggregated industry-wide data for knowledge extraction. It is obvious that this approach will be significantly less effective than the approach of integrating the detailed data from all competing companies. It is expected that, if these obstacles can be overcome, the impact of the DM and KDD methods will be much more prominent in industrial and commercial applications.

## References

- 36.1 M. J. A. Berry, G. Linoff: *Mastering Data Mining: The Art and Science of Customer Relationship Management* (Wiley, New York 2000)
- 36.2 E. Wegman: Data Mining Tutorial, Short Course Notes, Interface 2001 Symposium, Cosa Mesa, Californien (2001)
- 36.3 P. Adriaans, D. Zantinge: *Data Mining* (Addison-Wesley, New York 1996)
- 36.4 J. H. Friedman: *Data Mining and Statistics: What is the Connection? Technical Report* (Stat. Dep., Stanford University 1997)
- 36.5 K. B. Clark, T. Fujimoto: Product Development and Competitiveness, *J. Jpn Int. Econ.* **6**(2), 101–143 (1992)
- 36.6 D. W. LaBahn, A. Ali, R. Krapfel: New Product Development Cycle Time. The Influence of Project and Process Factors in Small Manufacturing Companies, *J. Business Res.* **36**(2), 179–188 (1996)
- 36.7 J. Han, M. Kamber: *Data Mining: Concept and Techniques* (Morgan Kaufmann, San Francisco 2001)
- 36.8 T. Hastie, J. H. Friedman, R. Tibshirani: *Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer, Berlin Heidelberg New York 2001)
- 36.9 S. Weisberg: *Applied Linear Regression* (Wiley, New York 1980)
- 36.10 G. Seber: *Multivariate Observations* (Wiley, New York 1984)
- 36.11 J. Neter, M. H. Kutner, C. J. Nachtsheim, W. Wasserman: *Applied Linear Statistical Models*, 4th edn. (Irwin, Chicago 1996)
- 36.12 A. E. Hoerl, R. Kennard: Ridge Regression: Biased Estimation of Nonorthogonal Problems, *Technometrics* **12**, 55–67 (1970)
- 36.13 R. Tibshirani: Regression Shrinkage and Selection via the Lasso, *J. R. Stat. Soc. Series B* **58**, 267–288 (1996)
- 36.14 A. Agresti: *An Introduction to Categorical Data Analysis* (Wiley, New York 1996)
- 36.15 D. Hand: *Discrimination and Classification* (Wiley, Chichester 1981)

- 36.16 P. McCullagh, J. A. Nelder: *Generalized Linear Models*, 2nd edn. (Chapman Hall, New York 1989)
- 36.17 T. Hastie, R. Tibshirani: *Generalized Additive Models* (Chapman Hall, New York 1990)
- 36.18 W. S. Cleveland: Robust Locally-Weighted Regression and Smoothing Scatterplots, *J. Am. Stat. Assoc.* **74**, 829–836 (1979)
- 36.19 R. L. Eubank: *Spline Smoothing and Nonparametric Regression* (Marcel Dekker, New York 1988)
- 36.20 G. Wahba: *Spline Models for Observational Data*, Applied Mathematics, Vol. 59 (SIAM, Philadelphia 1990)
- 36.21 W. Härdle: *Applied Non-parametric Regression* (Cambridge Univ. Press, Cambridge 1990)
- 36.22 D. Biggs, B. deVillie, E. Suen: A Method of Choosing Multiway Partitions for Classification and Decision Trees, *J. Appl. Stat.* **18**(1), 49–62 (1991)
- 36.23 B. D. Ripley: *Pattern Recognition and Neural Networks* (Cambridge Univ. Press, Cambridge 1996)
- 36.24 L. Breiman, J. H. Friedman, R. A. Olshen, C. J. Stone: *Classification and Regression Trees* (Wadsworth, Belmont, California 1984)
- 36.25 J. N. Morgan, J. A. Sonquist: Problems in the Analysis of Survey Data, and a Proposal, *J. Am. Stat. Assoc.* **58**, 415–434 (1963)
- 36.26 A. Fielding: Binary segmentation: The Automatic Interaction Detector and Related Techniques for Exploring Data Structure. In: *The Analysis of Survey Data, Volume I: Exploring Data Structures*, ed. by C. A. O'Muircheartaigh, C. Payne (Wiley, New York 1977) pp. 221–258
- 36.27 W. Y. Loh, N. Vanichsetakul: Tree-Structured Classification Via Generalized Discriminant Analysis, *J. Am. Stat. Assoc.* **83**, 715–728 (1988)
- 36.28 W. D. Lo, Chaudhuri, W. Y. Loh, C. C. Yang: Generalized Regression Trees, *Stat. Sin.* **5**, 643–666 (1995)
- 36.29 W. Y. Loh, Y. S. Shih: Split-Selection Methods for Classification Trees, *Statistica Sinica* **7**, 815–840 (1997)
- 36.30 J. H. Friedman, T. Hastie, R. Tibshirani: Additive Logistic Regression: a Statistical View of Boosting, *Ann. Stat.* **28**, 337–407 (2000)
- 36.31 Y. Freund, R. Schapire: Experiments with a New Boosting Algorithm, *Machine Learning: Proceedings of the Thirteenth International Conference*, Bari, Italy 1996, ed. by M. Kaufmann, (Bari, Italy 1996) 148–156
- 36.32 L. Breiman: Bagging Predictors, *Machine Learning* **26**, 123–140 (1996)
- 36.33 J. H. Friedman: Greedy Function Approximation: a Gradient Boosting Machine, *Ann. Stat.* **29**, 1189–1232 (2001)
- 36.34 J. H. Friedman: Stochastic Gradient Boosting, *Computational Statistics and Data Analysis* **38**(4), 367–378 (2002)
- 36.35 J. H. Friedman: Multivariate Adaptive Regression Splines (with Discussion), *Ann. Stat.* **19**, 1–141 (1991)
- 36.36 J. H. Friedman, B. W. Silverman: Flexible Parsimonious Smoothing and Additive Modeling, *Technometrics* **31**, 3–39 (1989)
- 36.37 R. P. Lippmann: An Introduction to Computing with Neural Nets, *IEEE ASSP Magazine April*, 4–22 (1987)
- 36.38 S. S. Haykin: *Neural Networks: A Comprehensive Foundation*, 2nd edn. (Prentice Hall, Upper Saddle River 1999)
- 36.39 H. White: Learning in Neural Networks: a Statistical Perspective, *Neural Computation* **1**, 425–464 (1989)
- 36.40 A. R. Barron, R. L. Barron, E. J. Wegman: Statistical Learning Networks: A Unifying View, *Computer Science and Statistics: Proceedings of the 20th Symposium on the Interface 1992*, ed. by E. J. Wegman, D. T. Gantz, J. J. Miller (American Statistical Association, Alexandria, VA 1992) 192–203
- 36.41 B. Cheng, D. M. Titterton: Neural Networks: A Review from a Statistical Perspective (with discussion), *Stat. Sci.* **9**, 2–54 (1994)
- 36.42 D. Rumelhart, G. Hinton, R. Williams: Learning Internal Representations by Error Propagation. In: *Parallel Distributed Processing: Explorations in the Microstructures of Cognition, Vol. 1: Foundations*, ed. by D. E. Rumelhart, J. L. McClelland (MIT, Cambridge 1986) pp. 318–362
- 36.43 V. Vapnik: *The Nature of Statistical Learning* (Springer, Berlin Heidelberg New York 1996)
- 36.44 C. J. C. Burges: A Tutorial on Support Vector Machines for Pattern Recognition, *Knowledge Discovery and Data Mining* **2**(2), 121–167 (1998)
- 36.45 J. Shawe-Taylor, N. Cristianini: *Kernel Methods for Pattern Analysis* (Cambridge Univ. Press, Cambridge 2004)
- 36.46 N. Cristianini, J. Shawe-Taylor: *An Introduction to Support Vector Machines* (Cambridge Univ. Press, Cambridge 2000)
- 36.47 P. Huber: *Ann. Math. Stat.*, Robust Estimation of a Location Parameter **53**, 73–101 (1964)
- 36.48 B. V. Dasarthy: *Nearest Neighbor Pattern Classification Techniques* (IEEE Computer Society, New York 1991)
- 36.49 T. Hastie, R. Tibshirani: Discriminant Adaptive Nearest-Neighbor Classification, *IEEE Pattern Recognition and Machine Intelligence* **18**, 607–616 (1996)
- 36.50 T. Hastie, R. Tibshirani, A. Buja: Flexible Discriminant and Mixture Models. In: *Statistics and Artificial Neural Networks*, ed. by J. Kay, M. Titterton (Oxford Univ. Press, Oxford 1998)
- 36.51 J. R. Koza: *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (MIT, Cambridge 1992)
- 36.52 W. Banzhaf, P. Nordin, R. E. Keller, F. D. Francone: *Genetic Programming: An Introduction* (Morgan Kaufmann, San Francisco 1998)
- 36.53 P. W. H. Smith: Genetic Programming as a Data-Mining Tool. In: *Data Mining: A Heuristic Approach*,



- ed. by H. A. Abbass, R. A. Sarker, C. S. Newton (Idea Group Publishing, London 2002) pp.157–173
- 36.54 R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, A. I. Verkamo: *Fast Discovery Of Association Rules: Advances in Knowledge Discovery and Data Mining* (MIT, Cambridge 1995) Chap.12
- 36.55 A. Gordon: *Classification*, 2nd edn. (Chapman Hall, New York 1999)
- 36.56 J. A. Hartigan, M. A. Wong: A K-Means Clustering Algorithm, *Appl. Stat.* **28**, 100–108 (1979)
- 36.57 L. Kaufman, P. Rousseeuw: *Finding Groups in Data: An Introduction to Cluster Analysis* (Wiley, New York 1990)
- 36.58 M. Ester, H.-P. Kriegel, J. Sander, X. Xu: A Density-Based Algorithm for Discovering Cluster in Large Spatial Databases, *Proceedings of 1996 International Conference on Knowledge Discovery and Data Mining (KDD96)*, Portland 1996, ed. by E. Simoudis, J. Han, U. Fayyad (AAAI Press, Menlo Park 1996) 226–231
- 36.59 J. Sander, M. Ester, H.-P. Kriegel, X. Xu: Density-Based Clustering in Spatial Databases: The Algorithm DGBSCAN and its Applications, *Data Mining and Knowledge Discovery* **2(2)**, 169–194 (1998)
- 36.60 M. Ankerst, M. M. Breunig, H.-P. Kriegel, J. Sander: OPTICS: Ordering Points to Identify the Clustering Structure, *Proc. ACM SIGMOD Int. Conf. on Management of Data*, Philadelphia, Pennsylvania June 1–3, 1999 (ACM Press, New York 1999) 49–60
- 36.61 T. Kohonen: *Self-Organization and Associative Memory*, 3rd edn. (Springer, Berlin Heidelberg New York 1989)
- 36.62 D. Haughton, J. Deichmann, A. Eshghi, S. Sayek, N. Teebagy, H. Topi: A Review of Software Packages for Data Mining, *Amer. Stat.* **57(4)**, 290–309 (2003)
- 36.63 J. R. Quinlan: *C4.5: Programs for Machine Learning* (Morgan Kaufmann, San Mateo 1993)
- 36.64 T. Fawcett, F. Provost: Activity Monitoring: Noticing Interesting Changes in Behavior, *Proceedings of KDD-99*, San Diego 1999, (San Diego, CA 1999) 53–62
- 36.65 D. C. Montgomery: *Introduction to Statistical Quality Control*, 5th edn. (Wiley, New York 2001)
- 36.66 W. H. Woodall, K.-L. Tsui, G. R. Tucker: A Review of Statistical and Fuzzy Quality Control Based on Categorical Data, *Frontiers in Statistical Quality Control* **5**, 83–89 (1997)
- 36.67 D. C. Montgomery, W. H. Woodall: A Discussion on Statistically-Based Process Monitoring and Control, *J. Qual. Technol.* **29**, 121–162 (1997)
- 36.68 A. J. Hayter, K.-L. Tsui: Identification and Qualification in Multivariate Quality Control Problems, *J. Qual. Tech.* **26(3)**, 197–208 (1994)
- 36.69 R. L. Mason, C. W. Champ, N. D. Tracy, S. J. Wierda, J. C. Young: Assessment of Multivariate Process Control Techniques, *J. Qual. Technol.* **29**, 140–143 (1997)
- 36.70 W. Jiang, S.-T. Au, K.-L. Tsui: A Statistical Process Control Approach for Customer Activity Monitoring, Technical Report, AT&T Labs (2004)
- 36.71 M. West, J. Harrison: *Bayesian Forecasting and Dynamic Models*, 2nd edn. (Springer, New York 1997)
- 36.72 C. Fraley, A. E. Raftery: Model-based Clustering, Discriminant Analysis, and Density Estimation, *J. Amer. Stat. Assoc.* **97**, 611–631 (2002)
- 36.73 G. Taguchi: *Introduction to Quality Engineering: Designing Quality into Products and Processes* (Asian Productivity Organization, Tokyo 1986)
- 36.74 G. E. P. Box, R. N. Kacker, V. N. Nair, M. S. Phadke, A. C. Shoemaker, C. F. Wu: Quality Practices in Japan, *Qual. Progress* **March**, 21–29 (1988)
- 36.75 V. N. Nair: Taguchi's Parameter Design: A Panel Discussion, *Technometrics* **34**, 127–161 (1992)
- 36.76 K.-L. Tsui: An Overview of Taguchi Method and Newly Developed Statistical Methods for Robust Design, *IIE Trans.* **24**, 44–57 (1992)
- 36.77 K.-L. Tsui: A Critical Look at Taguchi's Modeling Approach for Robust Design, *J. Appl. Stat.* **23**, 81–95 (1996)
- 36.78 G. Taguchi, S. Chowdhury, Y. Wu: *The Mahalanobis-Taguchi System* (McGraw-Hill, New York 2001)
- 36.79 G. Taguchi, R. Jugulum: *The Mahalanobis-Taguchi Strategy: A Pattern Technology System* (Wiley, New York 2002)
- 36.80 W. H. Woodall, R. Koudelik, K.-L. Tsui, S. B. Kim, Z. G. Stoumbos, C. P. Carvounis: A Review and Analysis of the Mahalanobis-Taguchi System, *Technometrics* **45(1)**, 1–15 (2003)
- 36.81 A. Kusiak, C. Kurasek: Data Mining of Printed-Circuit Board Defects, *IEEE Transactions on Robotics and Automation* **17(2)**, 191–196 (2001)
- 36.82 A. Kusiak: Rough Set Theory: A Data Mining Tool for Semiconductor Manufacturing, *IEEE Transactions on Electronics Packaging Manufacturing* **24(1)**, 44–50 (2001)
- 36.83 A. Ultsch: *Information and Classification: Concepts, Methods and Applications* (Springer, Berlin Heidelberg New York 1993)
- 36.84 A. Y. Wong: A Statistical Approach to Identify Semiconductor Process Equipment Related Yield Problems, *IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, Paris 1997 (IEEE Computer Society, Paris, France 1997) 20–22
- 36.85 ANSI (2002). Am. Nat. Standards Institute, IPC-9261, In-Process DPMO and Estimated Yield for PWB
- 36.86 M. Baron, C. K. Lakshminarayan, Z. Chen: Markov Random Fields In Pattern Recognition For Semiconductor Manufacturing, *Technometrics* **43**, 66–72 (2001)
- 36.87 G. King: Event Count Models for International Relations: Generalizations and Applications, *International Studies Quarterly* **33(2)**, 123–147 (1989)
- 36.88 P. Smyth: Hidden Markov models for fault detection in dynamic systems, *Pattern Recognition* **27(1)**, 149–164 (1994)