

Clinical prediction rules

Simon T Adams,¹ Stephen H Leveson²

¹York Hospital, York YO31 8HE, UK

²Hull-York Medical School, Learning and Research Centre, York Hospital
Correspondence to: S Adams
rpbgt@hotmail.com

Accepted: 3 October 2011

Cite this as: *BMJ* 2012;344:d8312
doi: 10.1136/bmj.d8312

Clinical prediction rules are mathematical tools that are intended to guide clinicians in their everyday decision making. The popularity of such rules has increased greatly over the past few years. This article outlines the concepts underlying their development and the pros and cons of their use

In many ways much of the art of medicine boils down to playing the percentages and predicting outcomes. For example, when clinicians take a history from a patient they ask the questions that they think are the most likely to provide them with the information they need to make a diagnosis. They might then order the tests that they think are the most likely to support or refute their various differential diagnoses. With each new piece of the puzzle some hypotheses will become more likely and others less likely. At the end of the process the clinician will decide which treatment is likely to result in the most favourable outcome for the patient, based on the information they have obtained.

Given that the above process is the underlying principle of clinical practice, and bearing in mind the ever increasing time constraints imposed on people, it is unsurprising that a great deal of work has been done to help clinicians and patients make decisions. This work is referred to by many names: prediction rules, probability assessments, prediction models, decision rules, risk scores, etc. All describe the combination of multiple predictors, such as patient characteristics and investigation results, to estimate the probability of a certain outcome or to identify which intervention is most likely to be effective.^{1 2} Predictors are identified by “data mining”—the process of selecting, exploring, and modelling large amounts of data in order to discover unknown patterns or relations.³

Ideally, a reliable predictive factor or model would combine both a high sensitivity with a high specificity.^{4 5} In other words it would correctly identify as high a proportion as possible of the patients fated to have the outcome in question (sensitivity) while excluding those who will not have the outcome (specificity).⁶ In the table sensitivity can be defined as $A \div (A+C)$ and specificity as $D \div (B+D)$.

A good predictive factor is not the same as a strong risk factor.⁴ The positive predictive value of a predictive factor or model refers to its accuracy in terms of the proportion of




Fig 1 | Receiver operating characteristic (ROC) curve

patients correctly predicted to have the outcome in question ($A \div (A+B)$ in the table).⁷ A risk factor can be identified by calculating the relative risk (or odds ratio) of an outcome in patients with the factor in question compared with patients without it.⁴ If, however, the factor identified or the outcome being used is uncommon, it is of little clinical use as a predictive factor.^{4 7}

A good predictive factor or model shows a good fit between the probabilities calculated from the model and the outcomes actually observed, while also accurately discriminating between patients with and without the outcome.^{4 5} For example, if all patients with a measured observation of ≥ 0.5 die and all patients with the measured observation < 0.5 survive then the observed factor is a perfect predictor of survival.

Unfortunately, as a general rule sensitivity and specificity are mutually exclusive—as one rises the other falls. Since both are important to the development of predictive models, receiver operating characteristic (ROC) curves are used to visualise the trade-off between the two and express the overall accuracy of the model (fig 1).^{4 8 9} Sensitivity (true positive) is plotted on the y axis and 1-specificity (false positive) is plotted on the x axis.^{4 9} The closer a point is to the top left of the graph then the higher the area under the curve and the more accurate or useful a predictive factor can be said to be.^{4 8 9} Conversely a plot in the 45 degree diagonal (denoting an area under the curve of 50%) indicates a test no more accurate than chance.^{4 8 9} Where the limits of acceptability are set is arbitrary and depends on several factors such as the severity of the outcome and the potential negative consequences of the test.^{4 9}

Establishing a clinical prediction rule

The establishment of a prediction model in clinical practice requires four distinct phases:

Tabular representation of predicted versus actual outcomes of a predictive model		
Predicted outcome	Actual outcome	
	Positive	Negative
Positive	A (true positive)	B (false positive)
Negative	C (false negative)	D (true negative)

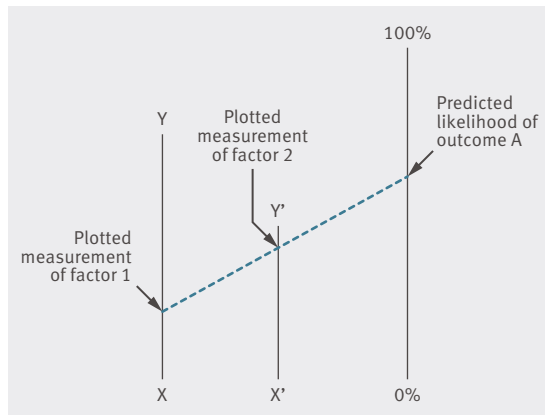


Fig 2 | Simplified representation of a basic nomogram

Development—Identification of predictors from an observational study

Validation—Testing of the rule in a separate population to see if it remains reliable

Impact analysis—Measurement of the usefulness of the rule in the clinical setting in terms of cost-benefit, patient satisfaction, time/resource allocation, etc

Implementation—Widespread acceptance and adoption of the rule in clinical practice.

For a prediction rule to gain popularity each of the first three steps needs to be satisfactorily completed before the fourth stage.¹ Validation in a suitably powered cohort study or controlled trial is particularly important because there is no guarantee that a predictor will be accurate outside the original dataset.¹⁻² Indeed validation usually shows a reduction in accuracy compared to that in the original study.¹⁻¹⁰⁻¹² Reliability is essentially the reproducibility of a measurement—that is, if the same test were applied under the same circumstances how similar the results would be.

Despite the long running controversy concerning their usefulness and application, the popularity of clinical prediction rules has been shown to be greater now than ever.¹⁻¹³⁻¹⁴ A Medline search by Toll and colleagues in 2008 showed that the number of papers discussing prediction rules has more than doubled in recent years (6744 papers in 1995 versus 15 662 in 2005).¹ Most publications, however, concern the development of new rules, with few articles describing validation and almost none confirming their clinical impact.¹ There are several possible reasons why validation and impact analysis are so often overlooked. Perhaps the most important are that neither validity nor reliability can be exactly quantified and that establishing validity requires investigators to consider several different aspects (face validity, content validity, construct validity, criterion validity, etc).¹⁵⁻¹⁶

Advantages and disadvantages of prediction rules

When appropriately developed and validated, prediction models have inherent advantages over human clinical decision making. Firstly, the statistical models can accommodate many more factors than the human brain is capable of taking into consideration.¹⁷ Secondly, if given identical data a statistical model will always give the same result whereas human clinical judgment has been shown to result in both inconsistency and disparity, especially with less experienced clinicians.¹⁷⁻¹⁸ Finally, and perhaps most importantly, several

prediction models have been shown to be more accurate than clinical judgment alone.¹⁴⁻¹⁷⁻²¹ So why are such models not used more readily in every practice?

Liao and Mark proposed in 2003 that resistance to adopting prediction models may reflect tacit acknowledgment that clinicians do not know how to take advantage of such tools.¹⁷ They also suggested that such tools may not be thought user friendly and may not take into account the continual, dynamic way in which humans gather clinical information.¹⁷ Their final reason for low implementation of clinical prediction rules is the sheer number of models available.¹⁷ If multiple prediction rules exist for the same problem identifying the best one is difficult. Not only is it potentially very time consuming but differences in the methods used in the studies on which they are based may make reliable comparison impossible.¹¹⁻²² Part of the reason for the large number of prediction rules may be the wide variety of ways in which such tools can be developed.

Types of prediction model

In 2006 Grobman and Stamilio described five main methods used to develop clinical prediction models: scoring systems derived from univariate analysis, prediction models based on multivariate analysis, nomograms, artificial neural networks, and decision trees.

Scoring systems derived from univariate analysis

Factors shown to be significantly related to the outcome in observational studies are allocated a score or “weight.” The cumulative final score of all the risk factors present in a patient is used as an indicator of the likelihood of the outcome occurring.⁴ Well known examples of this type of prediction model include the Alvarado score for acute appendicitis and the modified Glasgow score for acute pancreatitis.²³⁻²⁴ These models are simple to devise and use but their accuracy is affected by the potential inclusion of non-independent risk factors and the arbitrary manner in which factors are weighted.⁴

Prediction models based on multivariate analysis

These are developed in a similar manner to the above scoring systems except that the analysis of the results from the observational study is more refined and therefore less likely to include any non-independent factors. The models typically use logistic regression analysis, which has the added advantage of expressing the relation between the predictive factors and the outcome in the form of odds ratios (the probability of an outcome occurring versus the probability that it will not).⁴ These are relatively easy to interpret and can also be used to assign weights in a less arbitrary fashion than in univariate models.⁴⁻²⁵ Nevertheless, multivariate analysis techniques are not completely reliable in eliminating bias from interaction of independent variables.⁴ Models using logistic regression are often well suited to being represented as a nomogram (see below).³

Nomograms

Nomograms are graphical calculating devices that represent mathematical relations or laws and allow the user to rapidly calculate complicated formulas to a practical precision (fig 2).²⁶ Nomograms may be as simple as the markings on a thermometer or more complex, such as the Sigggaard-Andersen

chart used to diagnose acid-base blood disorders.²⁷ The mathematics and statistics used to develop a nomogram can be equally simplistic or intricate.⁴ The advantage of nomograms is that the final prediction tool created is generally comparatively simple to use and in some cases more accurate than other prediction models for the same clinical problem.^{4 28} Other nomograms in common clinical use include those used to predict the likelihood of a patient having prostate cancer from their clinical examination and prostate specific antigen levels and those used to predict the peak expiratory flow rate of asthmatic patients based on their age and height.^{29 30}

Prediction using artificial neural networks

Artificial neural networks are mathematical or computational models based on the operation of biological neural networks.³¹ In biology, a nerve cell (or neurone) will receive input from numerous other nerve cells. It will then process all of the input it receives and either send off an action potential or not. Because these nerve cells are all interconnected they are referred to as networks. Artificial neural networks function along similar lines: multiple sources of information (input) are fed into the software program, which interprets it and produces a dichotomous output (fig 3). The main advantage of neural networks is that they can “learn” mathematical relations between a series of input variables and the corresponding output.³²⁻³⁵ This is achieved by inputting a set of data containing both the input data (the predictor variables) as well as the outcomes.^{32 33} With each new dataset entered the neural network is able to adjust the internal weights of the various pieces of input data and calculate the probability of a specific outcome.³²

Neural networks require little formal statistical training to develop and can implicitly detect complex non-linear relations between independent and dependent variables as well as all possible interactions between predictor variables.^{32 33} However, they have a limited ability to explicitly identify pos-




Fig 3 | Schematic representation of an artificial neural network. The first column (input layer) represents a piece of data that can be put in to the neural network programme. The circles in the second column (hidden layer) represent the neural network programme assigned weight or numerical significance of each piece of data entered in the input layer. The final column (output layer) represents the dichotomous predicted outcome for the information entered




Fig 4 | Simplified representation of a basic decision tree

sible causal relations, they are hard to use at the bedside, and they require greater computational resources than other prediction models.^{32 33} They are also prone to “overfitting”—when too many datasets are used in training the network causing it to effectively memorise the noise (irrelevant data) and reducing its accuracy.^{32 33} A final drawback to neural networks is that the development model is empirical and because it is a new technique methodological problems remain.³² In a direct comparison between neural networks and logistic regression models Tu and colleagues concluded that neural networks were better for predicting outcomes but that logistic regression was preferable when looking for possible causal relations between independent and dependent variables or when trying to understand the effect of predictor variables on an outcome.³²

Decision trees (CART analysis)

Classification and regression tree (CART) analysis uses non-parametric tests to evaluate data and progressively divide it into subgroups based on the predictive independent variables.⁴ The variables and discriminatory values used and the order in which the splitting occurs are produced by the underlying mathematical algorithm and are calculated to maximise the resulting predictive accuracy.⁴ CART analysis produces “decision trees,” which are generally easily understood and consequently translate well into everyday clinical practice (fig 4). By following the arrows indicated by the answers to each of the questions in the boxes clinicians will be directed to the predicted outcome for the patient. Examples of CARTs used in clinical practice include those to predict large oesophageal varices in cirrhotic patients and to predict the likelihood of hospital admission in patients with asthma.^{36 37} However, the CART model of prediction can be significantly less accurate than other models.^{28 38} This may be because the “leaves” on the trees contain too few data to be able to predict outcomes reliably.³

Conclusion

Each of the five main models has advantages and disadvantages, and no single model of prediction has been clearly shown to be superior to the others in all applications. As pressure on their time increases, doctors will need to become familiar with decision making tools and the statistical principles underlying them.

Contributors: STA wrote the original manuscript and subsequent revisions. He is the guarantor. SHL provided critical evaluation of the original manuscript, suggested revisions, and gave final approval for submission of the paper for consideration for publication.

Competing interests: None declared.

Provenance and peer review: Not commissioned; externally peer reviewed. References are in the version on bmj.com.