

Paper 124-27

Communicating the Results of Predictive Models to Non-Technical Audiences



Andrew H. Karp
Sierra Information Services, Inc.
Sonoma, California USA



Abstract

Many statisticians, business analysts and other quantitative analysts are required to present the results of their modeling efforts to non-quantitatively trained decision-makers, or to decision-makers whose level of statistical training may be very limited. In the predictive modeling realm, where very sophisticated statistical tools such as logistic regression, neural networks and Chi-Square Automatic Interaction Detection (CHAID) are commonly employed, the knowledge gap between the provider and consumer of the analysis can become quite wide. When this occurs, the fruits of the analysis may go un- or under-utilized.

This paper discusses several ways that the results of predictive models can be effectively communicated to non-technical decision-makers. No matter how elegant your model is, or how well it satisfies various statistical criteria, its usefulness within the organization is hampered unless those who are tasked with implementing it can understand what the model will do for them, and what the anticipated benefits are of deploying the model.

Some general observations

In my experience, many "end-users" or "consumers" of predictive models have many attributes that make them very valuable to their organization. A physician, for example, may be a highly regarded expert in a particular class of diseases and/or very well respected in the community. The vice-president of marketing may have years of industry and product knowledge, excellent leadership and sales skills, and in-depth understanding of the marketplace within which the products or services are sold. But, their level of quantitative training may be very

limited, which in turn limits their ability to understand what the data analyst is presenting them.

Data analysts therefore often have trouble communicating with the end-users of their craft. The Physician-in-Chief or Vice President of Marketing probably knows very little about maximum likelihood estimation, and probably prefers to not have it explained to them. Nor are they going to "buy in" to a model because of the obtained value of the Akaike Information Criterion (AIC) or that the $-2\text{LOG}L$ measure reported by, for example, PROC LOGISTIC, is "statistically significant."

In many situations the decision-maker to whom the results of the model will be explained may have extensive financial or other resources "at risk" based on their decision to deploy a model. While they may intuitively understand the value of a predictive model, they are reluctant to commit scarce financial or other organizational resources to a project just because a model is "statistically significant." Here are two scenarios:

Scenario One

Data analysts at a Health Maintenance Organization (HMO) are tasked to develop a model to predict which patients are more (or less) at risk for influenza. The results of their model will be used to mount a mail campaign to encourage those at (statistical) risk for influenza to receive a vaccination before the onset of flu season. If the model performs well, it will correctly "rule in" as many of the high-risk patients and "rule out" those with low risk for influenza. Properly deployed, the a predictive model will help the HMO reduce the incidence of influenza (and related

treatment costs) among its members while not vaccinating those for whom the risk is low.

Scenario Two

A bank is interested in reducing the incidence of customers leaving their institution for another. This situation, often called "churn" or "attrition," is very often very expensive for banks and other consumer entities such as telecommunication companies because the cost to acquire a new customer is often much higher than the cost to keep a current customer. Senior bank management is willing to invest significant financial resources, as well as "cultural change management" resources, to orient large parts of the organization towards understanding what attributes make a customer more or less likely to "quit." Based on the results of such a predictive model, executive compensation schemes, corporate branding, customer service policies and other core "processes" may be drastically changed.

In both scenarios, the data analyst has a two-fold mission. The first is to develop a model that has good statistical properties and effective predictive "power." Using tools such as the SAS Enterprise Miner(tm) or SAS/STAT procedures such as PROC LOGISTIC, the analyst can create (depending on data availability) models that may have many desirable statistical properties. The models may, for example, exhibit high sensitivity and specificity, low values of measures such as the previously-mentioned AIC statistic, good overall measures of adequacy such as the SCORE or -2LOGL statistic. Testing the selected model on a validation or hold-out sample may increase the analyst's confidence in the worth of their model if it is deployed by the organization.

Now it's time to present the model to the decision-maker. Will the Physician-in-Chief at the HMO leap to her feet with joy that the model you've developed the lowest AIC value of all models you've looked at? Probably not. Nor is it reasonable to expect that the Senior Vice President of Marketing at a bank is going to agree to spend millions of dollars on

a customer retention campaign just because the LeGrange Multiplier (or SCORE) measure has a low p-value. Most likely not.

Tools to Communicate the Results of Predictive Models

Fortunately, there are parts of the output generated by SAS tools such as PROC LOGISTIC and Enterprise Miner software that can help you communicate the results of your models to non-technically oriented decision makers.

Parameter Estimates

The direction (positive or negative) of the model's parameter estimates can be used to explain the relationship between the values of the predictor (independent) variables and the outcome under analysis. Positive parameter estimates mean that larger values of the predictor variables are associated with higher values (or in a logistic regression model, higher probability of event outcome) of the response (dependent variable). Conversely, negative parameter estimates means that higher values of the predictor variables are associated with lower probability of event outcome (in the logistic regression case) or lower values of the dependent variable.

Odds Ratios

In my experience, one of the best ways to communicate the results of a predictive model to clients is the odds ratio. It is computed by from the values of the parameter estimates (and many SAS procedures, such as PROC LOGISTIC and PROC FREQ) generate them, either by default or when appropriate options are specified.

The odds ratio explains the relationship between changes in the values of an independent variable and the odds of event outcome. When a qualitative predictor variable (usually a binary or zero/one independent variable) is used, it is easy to say, for example, "patients with a pre-existing diagnosis of COPD (Chronic Obstructive Pulmonary Disease) are five times more likely to have influenza than patients without this

diagnosis," or that "customers who have a safe deposit box are twice as likely to respond to the insurance offer than those who don't.

When a continuous independent variable is used, the odds ratios might be very small, depending on how your data are prepared. For example, a one pound change in an adult patient's weight might be "statistically significant," but clinically irrelevant to health care providers. But a 10 or 20 pound change might be. Tools such as the UNITS option in PROC LOGISTIC can create customized odds ratios at user-specified values of the independent variable. Creating customized odds ratios is easy to do and can improve the "odds" your clients will be able to grasp how the model is working.

Sensitivity/Specificity

The model's results can also be articulated using the concepts of sensitivity and specificity. These measures are available when the CTABLE (classification table) option is specified in PROC LOGISTIC, or can be computed by hand from the output generated by other SAS procedures, such as PROC FREQ.

Sensitivity measures the ability of a model to correctly predict, or "rule in" the event of interest among those observations under analysis in which the event occurred. So, a model sensitivity of 90% means that if 100 observations in your data set HAD the event of interest, the model correctly predicted that 90 of them would have the event.

Specificity, on the other hand, measures the ability of a model to "rule out" the event in those observations under analysis in which the event DID NOT occur. If the model specificity is 90%, that means that out of 100 observations that did NOT have the event of interest, the model correctly predicted that 90 of them would not have the event.

These terms are also frequently used to articulate the efficacy of diagnostic medical tests and procedures. A diagnostic test that has 100% sensitivity and 50% specificity, for

example, can correctly "rule in" all patients who have the clinical condition of interest, and correctly "rules out" half of the patients who really don't have the condition.

Context drives which of the two measures is most important in an applied modeling context. It's almost impossible to come up with a model that has 100% sensitivity and 100% specificity. In many marketing situations the analysts and project managers want models that find (that is correctly identify) as many responders to a product offer as possible. For them, high sensitivity is often more important than high specificity. In many medical situations, however, a model which correctly rules out a diagnosis (i.e., specificity) may be more important. In this context, a model which can correctly avoid expensive (and perhaps painful), unnecessary medical procedures

Financial Considerations

One of the best ways to articulate the results of your model is to portray it in terms of its likely financial impacts. If you know, for example, the cost per customer contact and the likely revenue/profit to be obtained from each sale, you can easily calculate the potential revenue/profit expected from implementing your model. These amounts can be calculated within the Enterprise Miner, or by hand if other SAS System analytic tools are utilized in the modeling effort.

Gains/Lift Charts

These charts are available from the Enterprise Miner or can be created using a combination of Data Step programming and a BASE SAS Procedure called PROC RANK. A gains chart tells you how well model "captures" observations with the event of interest, and can quickly communicate how well the model can be expected to perform.

The core of a gains chart is the allocation (or assignment, or "binning") of each observation under analysis in to 10 (sometimes 20) groups, with the observations in the first group having the highest probability of event outcome, etc.

From there, the gains chart calculates how many responders (i.e., observations which had the event of interest) were found in each of the 10 groups. Cumulative numbers of observations, and (separately) numbers of observations with the event, are also included in the chart. Gains charts can be augmented with financial variables showing cost to contact (treat), expected revenue, profit/loss, etc.

A lift chart, on the other hand, displays the results of a predictive model in a different way than are shown in a gains chart. Both are useful in exploring and communicating the effectiveness of a predictive model.

Conclusion

Predictive modeling often offers an organization a way to train its resources effectively, thus boosting profits, reducing costs, eliminating errors, etc. While the SAS System provides an unparalleled suite of tools for predictive modeling projects, the value of the project lies outside the satisfaction of statistical criteria. Rather, the value of the project can only be realized when non-technically trained decision makers can understand the value of, and likely outcomes to be obtained from, the model. Tools and techniques such as those presented in this paper are easily available methods by which to (hopefully) bridge the communication gap that often limits the appropriate application of predictive models in many organizations.

References

Rud, Olivia Parr, *Data Mining Cookbook: Modeling Data for Marketing, Risk, and Customer Relationship Management*, Wiley (2001)

SAS Institute, Inc., *Data Mining Using the Enterprise Miner: A Case Study Approach*, SAS Institute, Inc. (2000)

Author contact

Andrew H. Karp
President
Sierra Information Services, Inc.
19229 Sonoma Highway PMB 264
Sonoma, California 94115 USA
707 996 7380
SierralInfo@AOL.COM
www.SierralInformation.com

Copyright

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the United States of America and other countries. (r) indicates USA registration. Other brand or product names are registered trademarks or trademarks of their respective companies.