



Classification of Breast Cancer Cells Using JMP®

Marie Gaudard, North Haven Group



Table of Contents

Abstract.....	1
1. Introduction	1
2. The Wisconsin Breast Cancer Diagnostic Data Set.....	1
3. Training, validation and test sets	2
4. Data visualization	6
4.1 One variable at a time	6
4.2 Two variables at a time.....	7
4.3 More than two variables at a time.....	10
4.4 Comparison of analysis data sets.....	11
5. Logistic model	13
5.1 A two-predictor model for visualization	13
5.2 A comprehensive logistic model for classification	14
6. Recursive partitioning	16
7. Neural net models.....	20
7.1 Background	20
7.2 First model – Neural Net 1	22
7.3 Second model – Neural Net 2.....	26
8. Comparison of classification models.....	30
9. Conclusion.....	31
References	32
Acknowledgments.....	32
Recommended Reading	32
Contact Information.....	33

Abstract

This paper illustrates some of the features of JMP that support classification and data mining. We will utilize the Wisconsin Breast Cancer Diagnostic Data Set, a set of data used to classify breast lumps as malignant or benign based on the values of 30 potential predictors, obtained by measuring the nuclei of fluid removed using a fine needle aspirate. We begin by illustrating some visualization techniques that help build an understanding of the data set. After partitioning our data into a training set, a validation set and a test set, we fit four models to the training data. These include a logistic model, a partition model and two neural net models. We then compare the performance of these four models on the validation data set to choose one. The test set is used to assess the performance of this final model.

1. Introduction

The purpose of this paper is to illustrate a number of features of JMP that support classification and data mining. We acknowledge at the outset that JMP is not intended to be a complete data mining package. However, it contains a number of modeling techniques that can be profitably used in data mining. It is our goal to show how a classification model can be fit relatively efficiently using JMP capabilities.

To illustrate these techniques, we will use a published data set, the Wisconsin Breast Cancer Diagnostic data, which is described in detail in the next section. For now, we only mention that the response of interest is whether a tumor is malignant or benign, and we will attempt to classify into these two categories using 30 potential predictors. After we describe the Wisconsin study and the data set, we will show ways to visualize this data, and then we will fit four classification models using three techniques: logistic modeling, recursive partitioning, and neural nets. We conclude with a comparison of the four classification models, choosing the best one based on performance on a validation subset of our data set.

2. The Wisconsin Breast Cancer Diagnostic Data Set

The Wisconsin Breast Cancer Data Set arises in connection with diagnosing breast tumors based on a fine needle aspirate (Mangasarian, OL, et al., 1994). In this study, a small-gauge needle is used to remove fluid directly from the lump or mass. The fluid is placed on a glass slide and is stained, so as to reveal the nuclei of the cells. A software program is used to determine the boundaries

of the nuclei. A typical image consists of 10 to 40 nuclei. The software computes 10 characteristics for each nucleus: radius, perimeter, area, texture, smoothness, compactness, number of concave regions and size of concavities (a concavity is an indentation in the cell nucleus), symmetry, and fractal dimension of the boundary (a higher value means a less regular contour). (For more detail on these characteristics, see Street, WN, et al., 1993.)

A set of 569 images was processed as described above. Since a typical image can contain from 10 to 40 nuclei, the data was summarized. For each variable, the Mean, Max and Standard Error of the Mean were computed. These are the 30 variables in our data set. The model developed by the researchers utilized a linear programming method that identified separating hyperplanes. Using all 569 observations, a classification accuracy of 97.3 percent was achieved. Even more remarkably, the next 131 cases that were analyzed were classified with 100 percent accuracy. Our final model utilized only three variables: Mean Texture, Max Area, and Max Smoothness.

We will utilize this data set in illustrating some of JMP software's capabilities in the area of classification analysis. The data set can be downloaded from the UCI Machine Learning Repository at [http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic)). JMP versions of this data set with embedded scripts that can be used to develop predictive models can be downloaded from <http://www.jmp.com/fileexchange>.

Breast Cancer Data Files:

BreastCancerClassification_Raw.jmp – This consists of the raw data and only two columns defined by us, namely the **Data Set Indicator** column (to ensure the same analysis data sets) and the **Random Unif** column on which the previous column is based. The user can utilize this data file to recreate, from scratch, the analyses in this paper.

BreastCancerClassification_Scripts.jmp – This data file contains scripts for most analyses, but the user must run these, save columns, and add formula columns to complete the analyses.

BreastCancerClassification.jmp – This data file contains the scripts that are in the file above as well as all of the added columns.

3. Training, validation and test sets

Our plan will be to construct models using three approaches: logistic regression, recursive partitioning and neural nets. These will be constructed using a portion of our data called the **training set**. Since bias (resulting from a tendency to overfit) may be built into these models as a result of being fit using the training data, we will assess these models relative to their performance on a hold-out portion of our data,

an independent data set called the **validation set**. The best model will be chosen, and, again, because choosing a model based on the validation set can also lead to overfitting, we will confirm the performance of the chosen model on another independent data set, called the **test set**. (See Bishop, CM, 1995, and Ripley, BD, 1996.)

Once again, if you wish to recreate our analysis on your own, use the data table **BreastCancerClassification_Raw.jmp**. If you want to do part of the work on your own, use **BreastCancerClassification_Scripts.jmp**. Otherwise, you will find all of the following work included in our working data table, **BreastCancerClassification.jmp**.

We begin by splitting our data set of 569 rows into three portions: a training set consisting of about 60 percent of the data, a validation set consisting of about 20 percent of the data, and a test set consisting of the remaining 20 percent of the data. We accomplish this in JMP by defining a new column, called **Random Unif**, containing the formula **Random Uniform()**. We then define another new column, called **Data Set Indicator**, using a formula (shown in Figure 1) that assigns rows to one of the three data sets based on the value of the random uniform value assigned in the column **Random Unif**. Note that we have hidden and excluded the column **Random Unif**, since it is not needed in what follows. To **Hide** and **Exclude** a column, right-click on the column name in the **Columns** panel, and choose **Hide** and **Exclude**. Two little icons will appear to the right of the column name to indicate that it is hidden and excluded.



Figure 1. Formula defining assignment to data sets

Now we need to make these three data set assignments convenient for use as we explore the data and fit models. We do this through the use of **row state variables**. A row state variable is a column whose contents are **row states**, namely attributes, such as **Exclude**, **Hide** and **Select**, that are applied to certain rows.

To create a row state variable, we start by defining a new column. In the **Column Information** window, under **Data Type**, choose **Row State**. Then click **OK**. This defines the new column as a row state column. We create three new columns in this fashion, calling them **Training Set**, **Validation Set** and **Test Set**.

We would like the column **Training Set**, when applied to the data table, to exclude and hide all rows that are not part of the training set. To accomplish this, we first apply this configuration to the data table, namely, we exclude and hide all rows that are not part of the training set. We go to **Rows > Row Selection > Select Where**, and select all rows where **Training Set Indicator does not equal Training Set**. We then choose **Rows > Exclude** and **Rows > Hide**. This inserts **Exclude** and **Hide** icons next to the selected rows. Finally, deselect the rows by left-clicking in the lower left triangular region in the upper left of the data grid, as shown in Figure 2. This deselects the rows, but retains the **Exclude** and **Hide** row states.

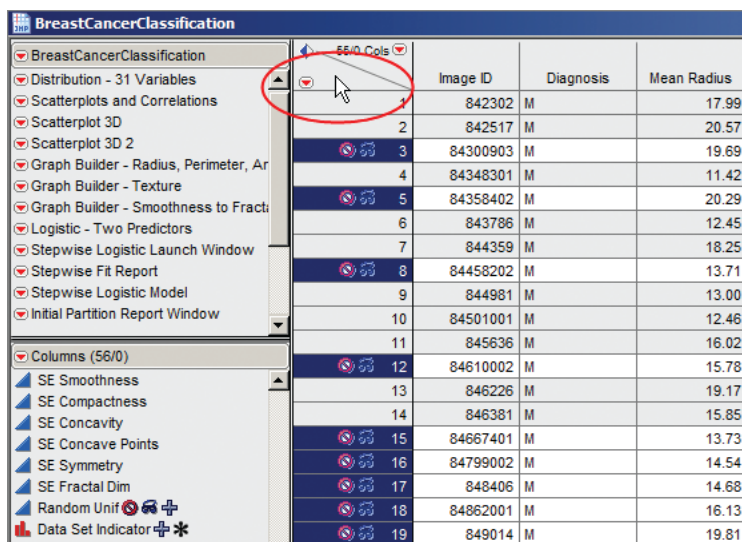


	Image ID	Diagnosis	Mean Radius
1	842302	M	17.99
2	842517	M	20.57
3	84300903	M	19.69
4	84348301	M	11.42
5	84358402	M	20.29
6	843786	M	12.45
7	844359	M	18.25
8	84458202	M	13.71
9	844981	M	13.00
10	84501001	M	12.46
11	845636	M	16.02
12	84610002	M	15.78
13	846226	M	19.17
14	846381	M	15.85
15	84667401	M	13.73
16	84799002	M	14.54
17	848406	M	14.68
18	84862001	M	16.13
19	849014	M	19.81

Figure 2. Lower triangular region of the data grid corresponding to rows

Finally, to insert these row states as values into the row state variable **Training Set**, we go to the **Columns** panel in the data table window, shown in Figure 3. Click on the star next to the row state variable, and choose **Copy from Row States**. We enter values into the three columns **Training Set**, **Validation Set** and **Test Set** using this procedure.

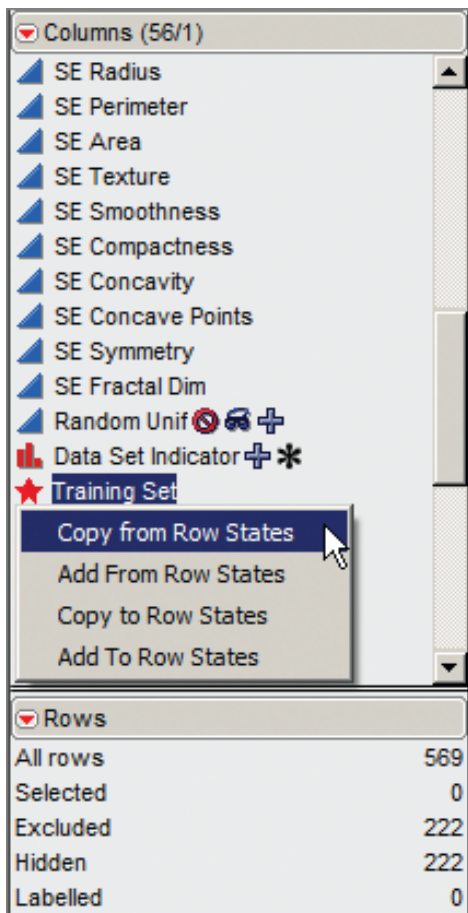


Figure 3. Columns panel showing choice for copying row states into the column Training Set

To further enhance our understanding of the data, it will be useful to color and/or mark the points in our plots according to the diagnosis. We will want to add this information to the row state variables that we have just created. Begin by going to **Rows > Clear Row States**. We will add colors and markers, and this will ensure that these are the only row states available at this time. Once row states are cleared, go to **Rows > Color or Mark by Column**, and click on **Diagnosis**. If you are following along on a computer, you can simply click **Make Legend with Window**, and JMP will color the points red and blue, based on whether the tumor is malignant or benign. A small legend window, which can be used to highlight points in plots, will also be created. To accommodate gray-scale printing of this paper, we also choose **Standard Markers** from the **Markers** drop-down menu. The markers for malignant masses are circles and the markers for benign masses are plus signs. When we click **OK**, the markers appear in the row margins. Now, for each row state variable in turn, we click the star to the left of that variable in the **Columns** panel and we choose **Add From Row States**. (Do this with care – it is easy to make mistakes!)

Just in case our colors and markers disappear, we create another row state column called **Diagnosis Colors**. We **Copy** from **Row States** so that this new variable contains the colors (and markers) that we have just defined. If you are creating these columns on your own, now is a good time to save your work!

To apply these row states as we need them, we will now simply be able to click on the star corresponding to the appropriate data set column, and choose **Copy to Row States**.

4. Data visualization

4.1 One variable at a time

At this point, no rows should be excluded or hidden. This can be verified by checking the Rows panel, shown in Figure 4, and noting that no rows are **Excluded** or **Hidden**. If there are excluded or hidden rows, go to **Rows > Clear Row States**. This removes all row states. Then go to the row state variable **Diagnosis Colors** in the **Columns** panel, and select **Copy to Row States** to reinstate the colors.

Rows	
All rows	569
Selected	0
Excluded	0
Hidden	0
Labelled	0

Figure 4. Rows panel

We want to see distribution reports for all of our variables. To do this, go to **Analyze > Distribution** and populate the launch window as shown in Figure 5.

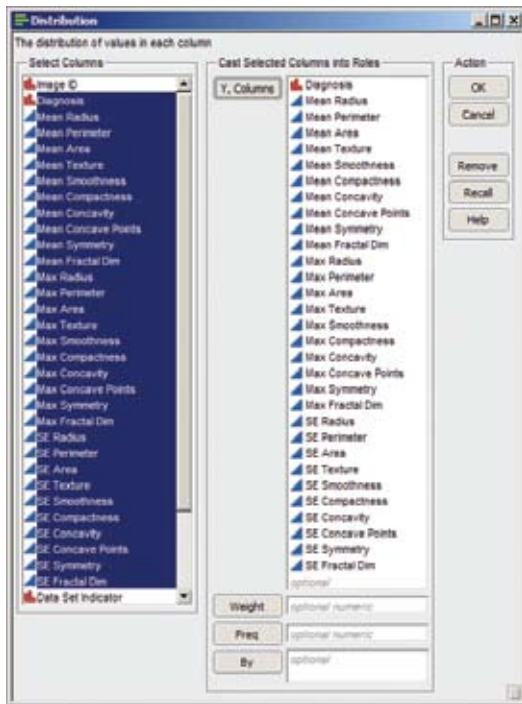


Figure 5. Launch window for distribution platform

Clicking **OK** results in 31 distribution reports, the first five of which are shown in Figure 6. The vertical layout for the graphs is the JMP default. However, this can be changed under **File > Preferences**. The bar graph corresponding to **Diagnosis** indicates that 212, or 37.258 percent, of the tumors included in the study were malignant. Scrolling through the plots for the 30 predictors, one can assess the shape of the distributions and the presence of outliers. We note that most distributions are skewed to the right and that there may be some outliers (for example, there may be two outliers for **SE Concavity**). One can also determine, by looking at **N** under **Moments**, that there are no missing values for any of the variables.

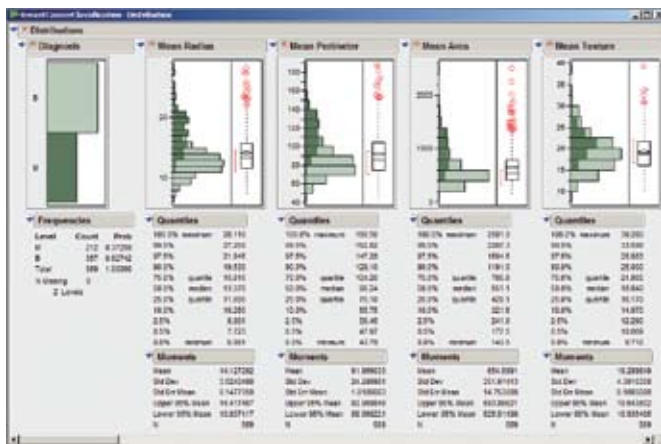


Figure 6. First five of 31 distribution reports

The script that generates the analysis in Figure 6 is saved to the data table as **Distribution – 31 Variables**. It can be found in the **Table** panel at the upper left of the data table. To run a script, simply click on the red triangle and choose **Run Script**.

4.2 Two variables at a time

In this section, we consider the issue of bivariate relationships among the variables. Of special interest is whether the predictors are useful in predicting **Diagnosis**.

For initial insight on this issue, we can utilize our **Distribution** output. Go to the bar graph for **Diagnosis**, and click on the bar corresponding to M. This has the effect of selecting all rows in the table for which **Diagnosis** has the value M. These rows are dynamically linked to all open plots, and so, in the 30 histograms corresponding to predictors, areas that correspond to the rows where **Diagnosis** is M are highlighted. We show five of the histograms corresponding to Max values in Figure 7. Note that masses that are malignant tend to have high values for these five variables. By scrolling through, one can detect relationships with most of the other variables as well. For example, malignant masses tend to have smaller **SE Texture** values than do benign masses. One can click on the bar for **Diagnosis** equal to B for additional insight.

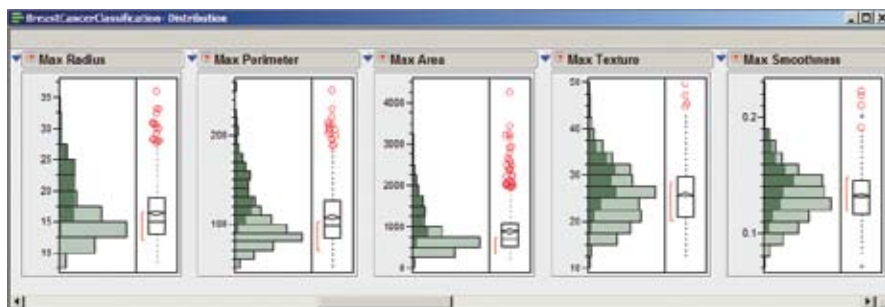


Figure 7. Five histograms for Max variables, with areas for malignant diagnoses highlighted

We are also interested in how the 30 predictors relate to each other. To see bivariate relationships among these 30 predictors, we will look at correlations and scatterplots. Go to **Analyze > Multivariate Methods > Multivariate**. In the launch window, enter all 30 predictors, from **Mean Radius** to **SE Fractal Dimension**, as **Y, Columns**. Clicking **OK** results in a report that shows a correlation matrix. Go to the red triangle and choose **Scatterplot Matrix** from the drop-down menu. This gives a 30 by 30 matrix showing all bivariate scatterplots for our 30 predictors (see Figure 8). The script that generates this output is saved to the data table as **Scatterplots and Correlations**.

We can think of radius, perimeter and area as variables that describe the size of the nuclei. From the 3-by-3 square at the top left of the scatterplot matrix (see Figure 8), we see that **Mean Radius** and **Mean Perimeter** are highly correlated (see

Correlations panel, which gives $r = 0.9979$), not an unexpected result. We also see that **Mean Area** is highly correlated with both **Mean Radius** and **Mean Perimeter**, with an expected quadratic relationship.

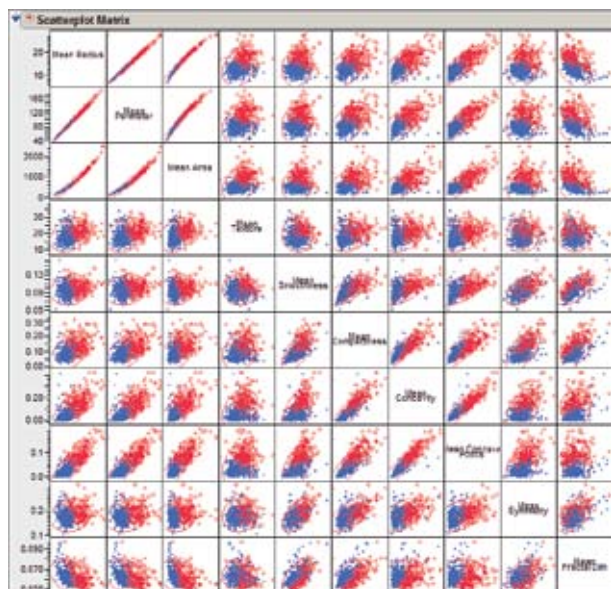


Figure 8. Portion of 30-by-30 scatterplot matrix corresponding to Mean variables

We note that the Max size variables are also highly correlated among themselves (all correlations greater than 0.9776), as are the SEs of the size variables (all correlations greater than 0.9377). These last details are easy to see if one clicks the red triangle next to **Multivariate**, and asks for **Pairwise Correlations**. (If you have run the script, this panel is already open.) Once in the **Pairwise Correlations** panel of the report, right-click, choose **Sort by Column**, and sort by **Correlation**.

It is also of interest to note that the Max size variables are fairly highly correlated with the Mean size variables. Again, this is not an unexpected result.

The effects of multicollinearity for explanatory models are well-known. For predictive models, such as the ones that we will be constructing, the impact of multicollinearity is minimal, so long as future observations come from within the range of the multicollinear relationships. In other words, if multicollinear relationships are present, we can interpolate, but we should not extrapolate.

Further examination of the scatterplot matrix suggests that there may be a few bivariate outliers. We choose not to pursue these at this point. However, one might want to consider their inclusion in model-building since they have the potential to be influential.

Incidentally, this is a nice opportunity to remind the reader of one of the shortfalls of interpreting correlations of grouped data. Consider the two variables **Mean Area** and **Mean Compactness**. From the correlation matrix, we see that their correlation is 0.4985, which seems somewhat substantial. To see the relationship between these two variables more clearly, we go to **Analyze > Multivariate Methods > Multivariate**,

and enter only **Mean Area** and **Mean Compactness** as **Y, Columns**. Clicking **OK** gives the output in Figure 9 (the script is **Scatterplot Matrix – Two Variables**). There does appear to be some correlation.

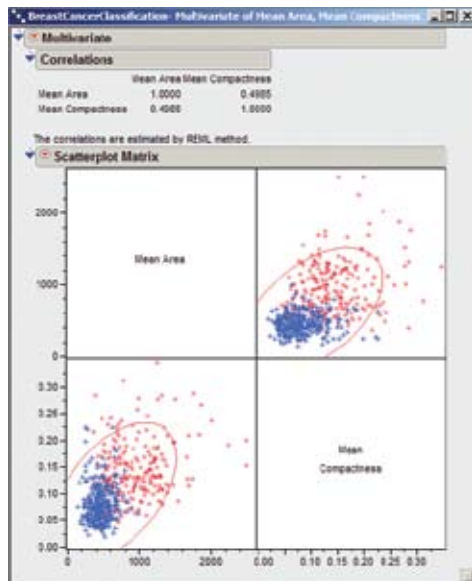


Figure 9. Correlations and Scatterplot Matrix for Mean Area and Mean Compactness

Now, we go back to **Analyze > Multivariate Methods > Multivariate**, click on **Recall** to repopulate the menu with the previous entries, and add **Diagnosis** as a By variable (the script is **Scatterplots – By Diagnosis**). The resulting output, shown in Figure 10, shows very little correlation between these two variables, based on **Diagnosis** grouping. The apparent correlation when the data is aggregated is a function of how the two **Diagnosis** groups differ relative to the magnitudes of the two predictors.

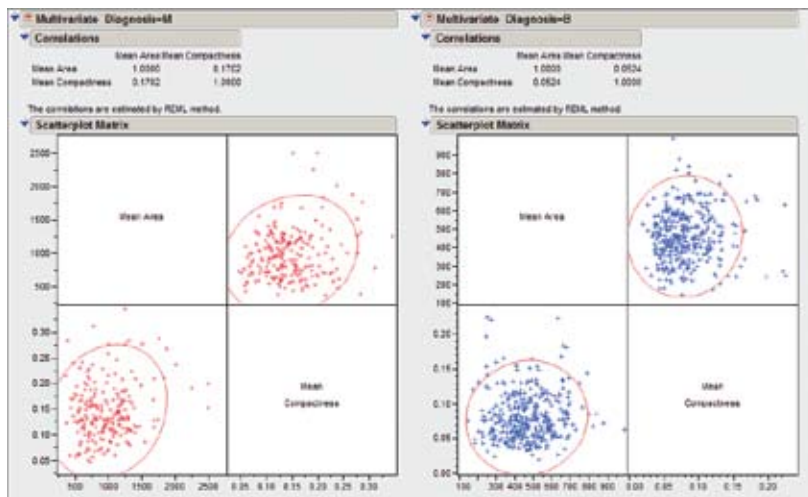


Figure 10. Correlations and scatterplots for Mean Area and Mean Compactness by diagnosis

4.3 More than two variables at a time

JMP provides several ways to visualize high-dimensional data. We will only look at a few of these.

To see the three-dimensional relationship among the three size variables, **Mean Radius**, **Mean Perimeter**, and **Mean Area**, we will go to **Graph > Scatterplot 3D**. Enter the three size variables as **Y, Columns**, and click **OK**. The resulting plot is a rotatable 3D plot of the data, shown in Figure 11. The script generating this plot is called **Scatterplot 3D**. The dependencies of the three variables are striking. (Note that the markers and colors and the **Legend** window must be previously constructed through the **Rows** menu – this is not captured in the **Scatterplot 3D** script.)

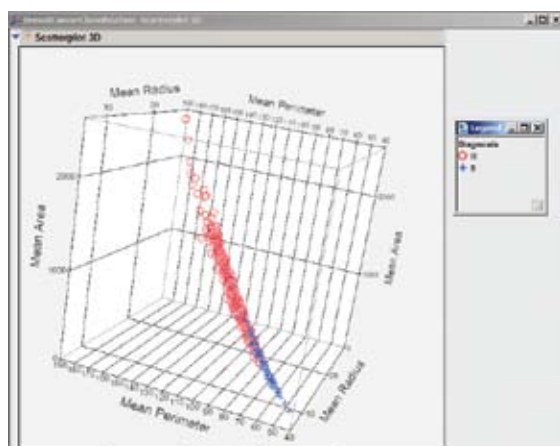


Figure 11. Three-dimensional view of Mean size variables, by diagnosis

It is possible to enter all 30 predictor variables as **Y, Columns** in the Scatterplot 3D launch window. When that is done, JMP allows the user to select three variables from the drop-down menus at the bottom of the plot to be used for the axes. Figure 12 displays the 3D scatterplot for three of the Max variables. Note that the two diagnosis classes have a reasonable degree of separation in this three-dimensional space. The script that generates this report is called **Scatterplot 3D 2**, and is saved to the data table.

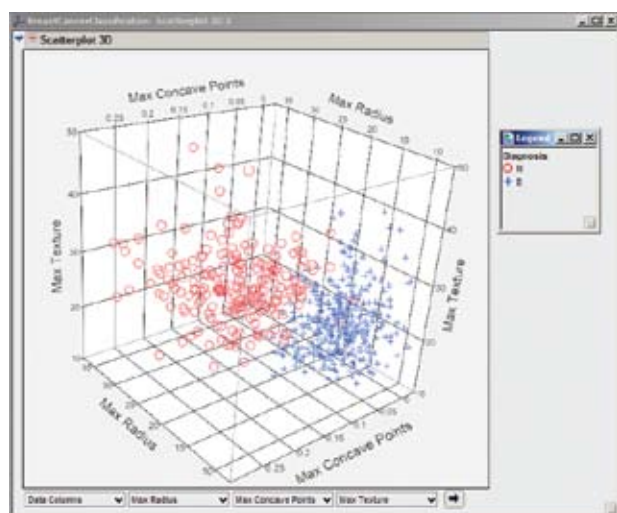


Figure 12. Scatterplot 3D for Max Radius, Max Concave Points and Max Texture, by diagnosis

4.4 Comparison of analysis data sets

We might be interested in confirming that our training, validation, and test sets do not differ greatly in terms of the distributions of the predictors. JMP software's **Graph Builder** (new in JMP 8), which is found under **Graph**, provides an intuitive interface for comparing the distributions of multiple variables in a variety of ways. The graph in Figure 13 utilizes boxplots to compare the distributions of the variable **Area**, for the **Mean**, **Max**, and **SE** groupings, across the training, validation and test sets. (The script that creates this output is saved as **Graph Builder – Area**.) We conclude that the distributions across the analysis sets are quite similar, with the exception of two outliers for **SE Area** that appear in the Validation and Test sets. We also note that the shapes of the distributions for the **Mean** and **Max** size values are similar as well.

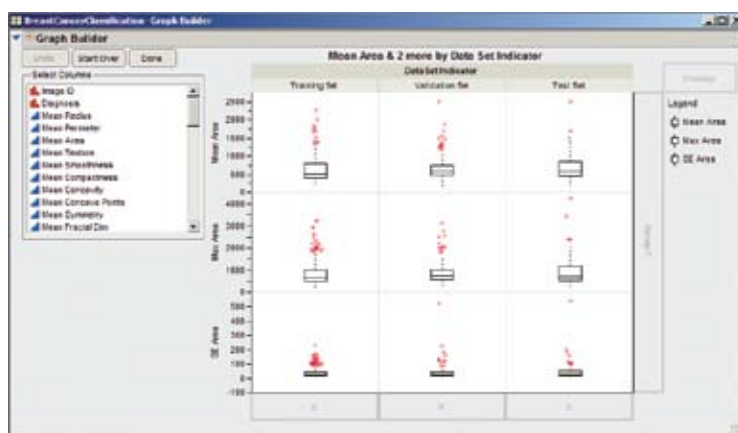


Figure 13. Boxplots of Mean, Max and SE of Area, by analysis data set

Figure 14 compares the distributions of **Radius**, **Perimeter** and **Texture** across analysis sets, using **Graph Builder**. The script that generates this graph is **Graph Builder – Radius, Perimeter, Texture**. Again, the distributions are quite similar, but we note the outliers on **SE Perimeter** for the validation and test sets. By selecting these two points in the graph in Figure 14, the points for the corresponding rows become selected in the graph in Figure 13, and we see that the outliers for **SE Perimeter** are the same rows as are outliers for **SE Area**.

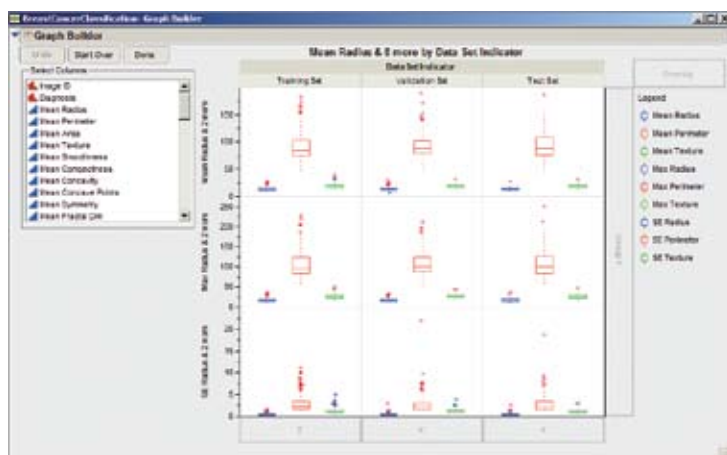


Figure 14. Boxplots of Mean, Max, and SE of Radius, Perimeter, and Texture, by analysis data set

Now, we could have obtained boxplots for all 10 variables in the format shown in Figure 14. But, because of the differences in scaling, the boxplots for the rest of the variables would have appeared as horizontal lines, showing no detail whatsoever. So, a caveat in using **Graph Builder** is that variables need to be comparable in terms of scale in order for the type of display shown in Figure 14 to be meaningful.

Boxplots for the remaining variables, **Smoothness**, **Compactness**, **Concavity**, **Concave Points**, **Symmetry** and **Fractal Dim**, are shown in Figure 15. The script for this collection of graphs is called **Graph Builder – Smoothness to Fractal Dim**. Again, the distributions are roughly comparable across analysis data sets, although there is evidence of two outliers on **SE Concavity** in the training set.




Figure 15. Boxplots for Mean, Max and SE of Smoothness through Fractal Dim, by analysis data set

5. Logistic model

At this point, we will begin the process of fitting classification models to our data. We will utilize only the training set for fitting models. So, we copy the values in the variable **Training Set** to the row states in our data table (click on the star to the right of **Training Set** in the **Columns** panel and select **Copy to Row States**). In the **Rows** panel, we should see that 222 rows are **Excluded** and **Hidden**.

5.1 A two-predictor model for visualization

Before fitting our “official” logistic model for classification, we want to take the opportunity to show what a logistic model, based on only two predictors, might look like. We will construct a surface plot for the predicted probability that a mass is malignant, based on a logistic model with only two predictors: **Mean Perimeter** and **Mean Smoothness**. The script for the model is **Logistic – Two Predictors**. For those following along using the data table **BreastCancerClassification.jmp**, first run

this script. Then, under the red tab in the **Nominal Logistic Fit** report, choose **Save Probability Formula**. This will save four new formula columns to your data table. In particular, we are interested in the predicted probability of malignancy, which is called **Prob[M] 2**. (The “2” is added because the data table already contains a column called **Prob[M]** from our comprehensive analysis, which we will run shortly.)

The surface plot in Figure 16 is obtained by going to **Graph > Surface Plot**, and adding **Prob[M] 2**, **Mean Perimeter** and **Mean Smoothness** as **Columns**. The surface represents the predicted probability of a malignant diagnosis (**Prob[M] 2**). The values of **Prob[M] 2** are plotted for the points in our training set; recall that these points are represented as red circles for malignant tumors and blue plus signs for benign tumors. Note that a grid has been inserted at **Prob[M] = 0.5**. If you were to classify observations with **Prob[M]** values above 0.5 as malignant and below 0.5 as benign, then you can see that the classification rule would correctly classify a fair number of observations (to be exact, 312 of the 347 observations are correctly classified).

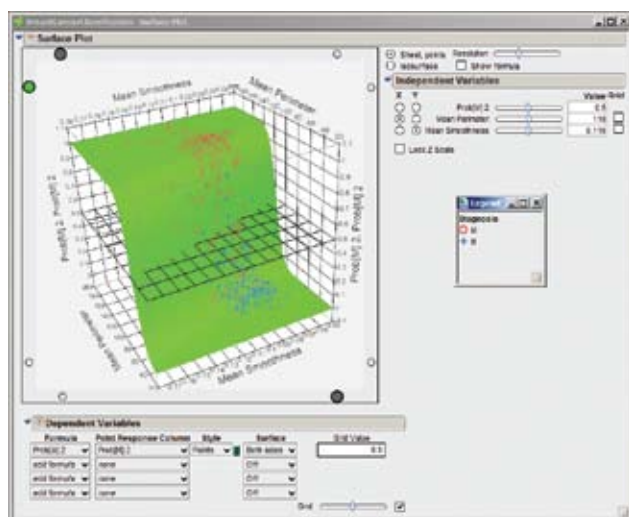


Figure 16. Logistic model based on Mean Perimeter and Mean Smoothness

It's important to realize that logistic regression results in an estimate of the probability of class membership, conditional on the values of the predictor variables. So, it makes sense to classify a new observation into the malignant class if and only if **Prob[M] 2** exceeds (or equals) 0.5. (If you are following along and have saved columns based on this analysis, please delete them now.)

5.2 A comprehensive logistic model for classification

At this point, we consider fitting a model based on all 30 potential predictors. Our first issue relates to variable selection. We are given 30 potential predictors, but we could expand this set if we wished. For example, two-way interactions (such as the interaction of **Mean Perimeter** and **Mean Texture**, which might be a useful predictor) and/or quadratic terms could be explored in fitting classification models. This brings us quickly to an intimate awareness of the well-known “curse of dimensionality.” Even with only 30 predictors, there are well over 1 billion different possible logistic models.

Why not simply use all 30 predictors? There are several reasons. Too many predictors can lead to overfitting. In other words, with enough predictors, it is possible to fit every point in a data set exactly. But such models tend not to generalize very well. So it is important to find a set of predictors that describes the structure of the data, and not the vagaries of individual points. Also, there can be computational issues when too many predictors are present. In particular, multicollinearity can degrade models obtained using certain algorithms (such as least squares regression).

For our breast cancer data, if we fit a logistic model to **Diagnosis** using all 30 predictors, we can perfectly classify every observation in our training set. But, when we apply the classification formula to our validation set, it behaves quite poorly. So it is important that we reduce the number of variables. We will do so using JMP software's stepwise variable selection procedure.

We go to **Analyze > Fit Model**, where we enter **Diagnosis** as **Y** and all 30 of our predictors in the **Construct Model Effects** box. Under **Personality**, we choose **Stepwise**. This window can be obtained by running the script **Stepwise Logistic Launch Window**.

Clicking **Run Model** opens the **Stepwise Fit** report. This window allows the user to make stepwise selection choices similar to those available in linear regression analysis. Just as a reminder, **forward** and **backward** selection procedures consist, respectively, of entering the most desirable variable and removing the least desirable variable at each step. We will do our variable selection using a **mixed** procedure, meaning that each **Enter** step will be followed by a **Remove** step. To communicate this to JMP, we set **Direction** equal to **Mixed**. We also set both the **Prob to Enter** and **Prob to Leave** values at 0.15. This will allow variables with significance level below 0.15 to enter, and variables that have entered at a previous stage, but which now have significance levels exceeding 0.15, will be available to be removed. The **Stepwise Regression Control** panel is shown in Figure 17; it can be obtained by running the script **Stepwise Fit Report**.

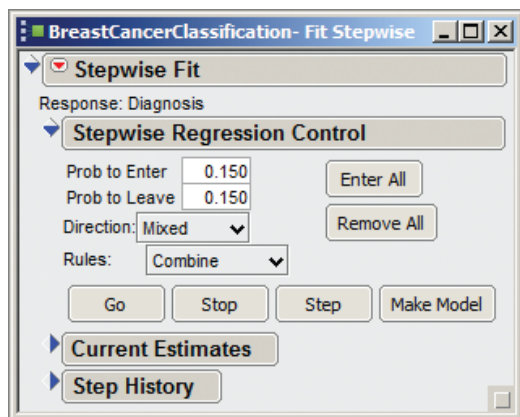


Figure 17. Stepwise Fit panel for logistic variable selection

When we click on **Go**, seven variables are selected. Then, we click on **Make Model**. This results in a **Fit Model** window containing the specification for our seven predictor model (the script is called **Stepwise Logistic Model**). The report that opens when we click **Run Model** shows details about the model fit. Under the red triangle, for example, we can ask to view the **Receiver Operating Characteristic (ROC)** curve or the **Lift** curve. We will ask to Save **Probability Formula**. This saves a number of formulas to the data table, including **Prob[M]** and **Prob[B]** – the probability that the tumor is malignant or benign, respectively – as well as a column called **Most Likely Diagnosis**, which gives the **Diagnosis** class with the highest probability, conditional on the values of the predictors.

To see just how well the classification performs on the training data, we can now go to **Analyze > Fit Y by X**. Enter **Most Likely Diagnosis** as **Y, Response**, and **Diagnosis** as **X, Factor**. The resulting mosaic plot and contingency table are shown in Figure 18. Of the 347 rows, only five are misclassified. We remind the reader, though, that there is inherent bias in evaluating a classification rule on the data used in fitting the model. In a following section, we will compare our three models using the validation data set.

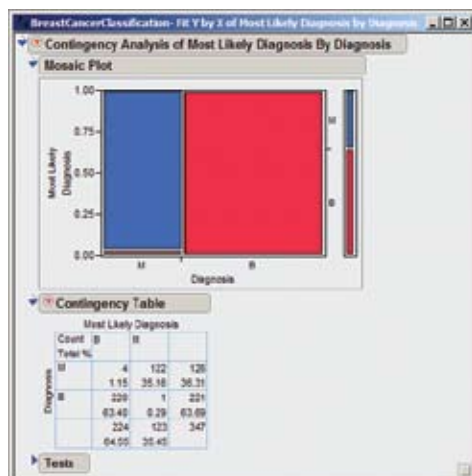


Figure 18. Mosaic Plot and Contingency Table for classification based on logistic model

6. Recursive partitioning

Our next model will be fit using JMP software's recursive partitioning platform, which provides a version of classification and regression tree analysis. The partition platform allows both response and predictors to be either continuous or categorical. Continuous predictors are split into two partitions according to cutting values, while categorical predictors (predictors that are nominal or ordinal) are split into two groups of levels.

If the response is continuous, the sum of squares due to the differences between means is a measure of the difference in the two groups. Both the variable to be split at a given level and the cutting value for the split are determined by maximizing a quantity, called the LogWorth, which is related to the p-value associated with the sum of squares due to the difference in means. In the case of a continuous response, the fitted values are the means within the two groups.

If the response is categorical, as in our case, the splits are determined by maximizing a LogWorth statistic that is related to the p-value of the likelihood ratio chi-square statistic, reported in the JMP output as G^2 . In this case, the fitted values are the estimated proportions, or response rates, within groups.

The partition platform is useful for both exploring relationships and for modeling. It is very flexible, allowing a user to find not only splits that are optimal in a global sense, but also node-specific splits that satisfy various criteria. The platform provides only a minimal **stopping rule** – that is, a criterion to end splitting. This rule is based on a user-defined minimum node size. The platform does not incorporate any other stopping rules; this is advantageous in that it enhances flexibility.

To fit our partition model, we go to **Analyze > Modeling > Partition**. Enter **Diagnosis** as **Y, Response**, and enter all 30 predictor variables as **X, Factor**. Clicking **OK** results in an initial report window. Under the red triangle, choose **Display Options > Show Split Prob** – this shows the split proportions in the nodes. This results in the report shown in Figure 19 (the script for this analysis is called **Initial Partition Report Window**). The graph shows the blue plus signs, indicating benign tumors, and the red circles, indicating malignant tumors, separated by a horizontal line at 0.36, the overall proportion malignant.

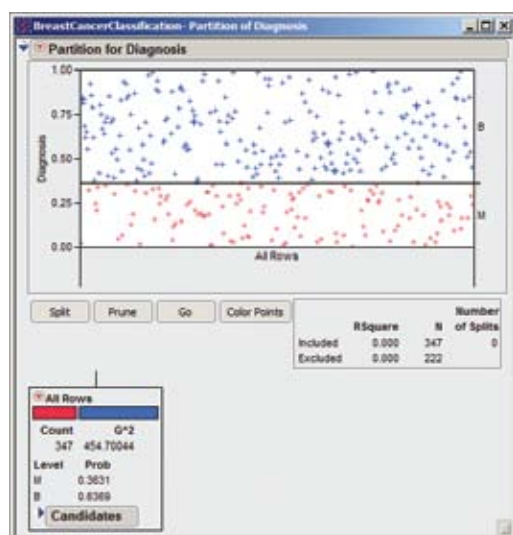


Figure 19. Initial partition report window

We begin to split by clicking on the **Split** button. JMP software's algorithm determines the best variable on which to split and the best cutting value of that variable. The tree, after the first split, is shown in Figure 20. The first split is on the variable **Max Concave Points**, and the observations are split at the value where **Max Concave Points** = 0.14. Of the observations where **Max Concave Points** \geq 0.14 (the rightmost node), 96.40 percent are malignant. Of those for which **Max Concave Points** $<$ 0.14, 91.95 percent are benign.

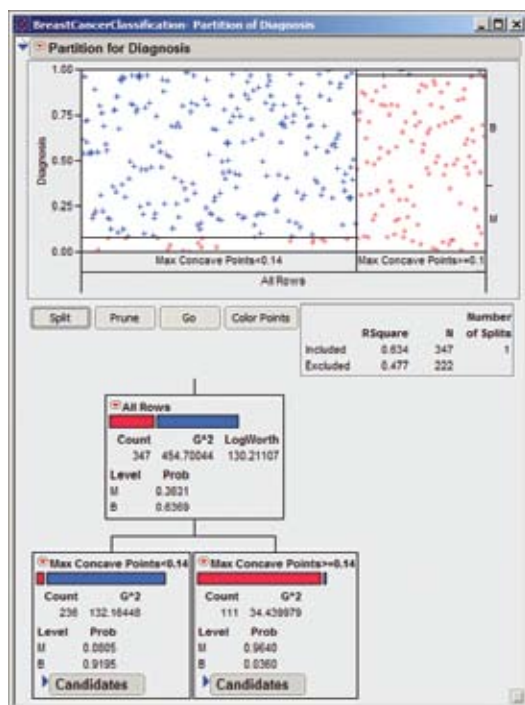


Figure 20. Partition report after first split

We continue to split for a total of eight splits. At this point, no further splits occur. This is because there is a default minimum split size of five and further splitting of nodes would result in nodes of smaller size than five. It is true that we could have stopped splitting before reaching eight splits. In fact, we may have overfit the data, given that some of the final nodes contain very few observations.

The red triangle at the top of the partition report gives the user many options, including **ROC Curve**, **Lift Curve**, **Leaf Report**, **K-Fold Cross Validation**, etc. The **Small Tree View** that is provided is shown in Figure 21 – it summarizes the structure of our splits. This report is given by the script **Partition Model**. Under the red triangle, we go to **Save Columns > Save Prediction Formula**. This saves two new columns to the data table, **Prob(Diagnosis==M)** and **Prob(Diagnosis==B)**. These columns give the predicted probabilities, which are simply terminal node proportions, of the respective **Diagnosis** classes. So, the classification rule consists of determining which node a new observation falls into, and predicting its probability of class membership based on the sample proportions in that node.

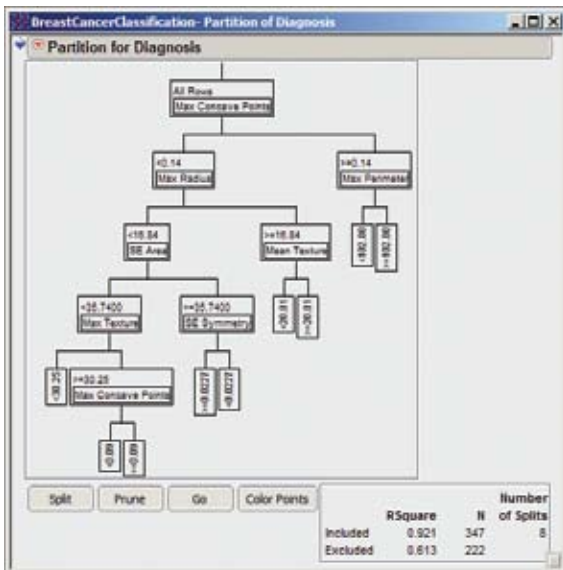


Figure 21. Small tree view for partition model

In the data table, we define a new column called **Partition Prediction**. We construct a formula to define this column; this formula classifies an observation as “M” if $\text{Prob}(\text{Diagnosis}=\text{M}) \geq 0.5$, and as “B” otherwise. Once this is done, we go to **Analyze > Fit Y by X**. We enter **Partition Prediction** as **Y, Response**, and **Diagnosis** as **X, Factor**. The resulting mosaic plot and contingency table are shown in Figure 22. Of the 347 rows, 11 are misclassified. At first blush, this model does not seem to be performing as well as the logistic model. In a later section, we will compare it to the logistic model by evaluating it on the validation data set.

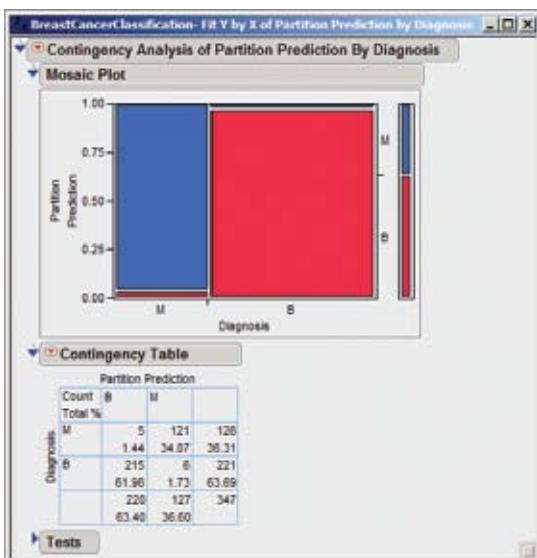


Figure 22. Mosaic Plot and Contingency Table for classification based on partition model

7. Neural net models

7.1 Background

As we mentioned earlier, data mining algorithms for classification and prediction can be based on **neural nets**. Neural net algorithms were originally inspired by how biological neurons are believed to function. Starting in the 1940s, scientists in the area of artificial intelligence pursued the idea of designing algorithms that “learn” in a fashion that emulates neurons in the human body.

The science of biologically informed computation had its origins in a seminal paper called “A Logical Calculus of Ideas Immanent in Nervous Activity,” by Warren McCulloch and Walter Pitts (both at MIT), published in 1943. This paper showed that the “ideas” in nerve cells were carried by the entire collection of neurons as a whole. They were implicit. (*Immanent*, in the title of the paper, means “having existence or effect only within the mind or consciousness” – Webster’s.)

Research since these early days has leveraged the idea of utilizing neuron-based processing elements arranged in nets to produce the algorithms that appear today in neural net software. Neural nets are based on neuron-like processing elements arranged in nets. There is an input layer of neurons, an output layer, and a hidden layer where processing occurs.

For some intuition on the way neural net algorithms work, consider the diagram in Figure 23, where a neural net is attempting to distinguish a golf ball from a football. Note that the neurons are arranged in three layers: input neurons, hidden neurons (the **hidden layer**), and output neurons. Neurons can be excited to various degrees – they are not completely “on” (excited, black) or “off” (unexcited, white). The input neurons are sensory, in that they react to a stimulus and output a value that reflects the intensity of that stimulus. A pattern of weights connects the remaining neurons and determines the degree to which each is excited. The output and hidden layer neurons sum their inputs and compare this sum to a threshold value to determine their output.

The net learns when it makes an error. The degree of error is propagated back through the network. Connections that were wrong are down-weighted, while those that were correct are strengthened. Active neurons with strong connections that were wrong are strongly penalized. In our golf ball and football example, prior to training, the connection strengths and excitation levels would be random.

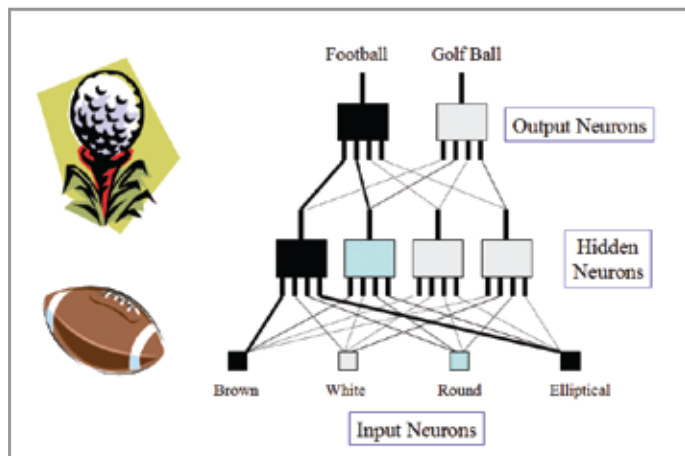


Figure 23. Schematic of a neural net being trained to distinguish a football from a golf ball

Now, in a mathematical sense, a neural net is nothing more than a nonlinear regression model. In its implementation of neural nets, JMP uses standard nonlinear least squares regression methods. Although a general neural net can have many hidden layers, one layer is considered sufficient for most modeling situations, and JMP utilizes a single hidden layer. Each hidden node is modeled using a logistic function applied to a linear function of the predictors. In a classification situation, such as ours, the output value consists of the logistic function applied to a linear function of the hidden nodes. This means that, for example, for 30 input variables, one response and k hidden nodes, the number of parameters to be estimated is $31 \cdot k + k + 1$.

With so many parameters, it is easy to see that a major advantage of a neural net is its ability to model a variety of response surfaces. But the large number of parameters comes at a cost. There are many local optima, and convergence to a global optimum can be difficult. Also, with so many parameters, overfitting is problematic. (This is why validation sets are critical to neural net modeling strategies.) Another disadvantage of neural net models is that they tend not to be interpretable, due to the hidden layer.

JMP software's implementation provides a user-specified overfitting penalty to help minimize overfitting issues. The overfitting penalty also helps mitigate the effects of multicollinearity. The user is able to set the number of nodes in the hidden layer. Here, a small number can lead to underfitting and a large number can lead to overfitting. Each application of the algorithm has a random start, and JMP refers to these individual fits as **tours**. About 20 tours are recommended in order to find a global optimum.

JMP also provides two methods to help a user select a neural net model that will extend well to new data. One of these methods is called **Random Holdback**. In this method, a sample of the observations is withheld (the **holdback** sample) and the remaining observations are used to train a neural net. JMP computes an R^2 value for the training sample, and then applies the neural net to the holdback sample and calculates an R^2 for that sample, which is called the **crossvalidation R^2** , denoted **CV**

RSquare in the JMP report. The user can vary the number of nodes and the overfit penalty in an attempt to find a model that generalizes well to the holdback sample. This method works well for large samples, where one can easily fit a model to 75 percent or fewer of the observations (JMP uses two-thirds of the complete data set by default).

The second method is called **K-Fold Crossvalidation**. Here, a neural net model is fit to all of the data to provide starting values. Then, the observations are divided randomly into K groups (or **folds**). Each of these groups, in turn, is treated as a holdback sample. For each of the K groups, a model is fit to the data in the other ($K - 1$) folds, using the starting values from the full fit. The model is then extended to the holdback group. An R^2 is calculated for each holdback sample, and these are averaged to give a **CV RSquare** that represents how the model might perform on new observations. (The starting values from the full fit are used because the function being optimized is multimodal, and this practice attempts to bias the estimates for the submodels to the mode of the overall fit.)

7.2 First model – Neural Net 1

We will begin by fitting a model to our breast cancer data without any use of crossvalidation. To fit a neural net model, we go to **Analyze > Modeling > Neural Net**. As usual, we enter **Diagnosis** as **Y, Response** and all 30 predictors as **X, Factors**. A diagram of a model with three hidden nodes is shown in Figure 24; this can be obtained by choosing **Diagram** under the red triangle in the report window. The final model will be a linear combination of three models, each of which relates one of the hidden nodes – H1, H2, and H3 – to the 30 predictors.

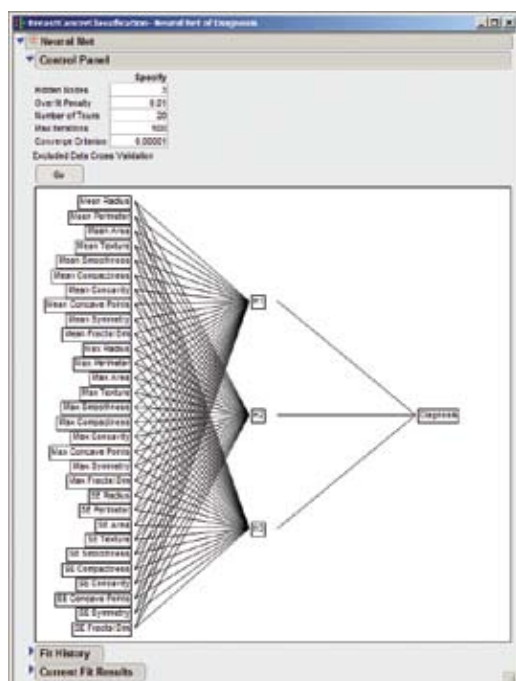


Figure 24. Diagram of neural net model for 30 predictors with three hidden nodes

The settings for the neural net are entered in the **Control Panel**, shown above the diagram in Figure 24 and repeated in Figure 25. Note that the overfit penalty in JMP 8 is set by default at 0.01, a value that we accept for this model. (Smaller values of the overfit penalty led to models with sharp curves, suggesting overfitting.) We accept these defaults and click **GO** to obtain the fitted model. Depending on the data set size and the settings in the **Control Panel**, the fitting procedure can take a noticeable amount of time.

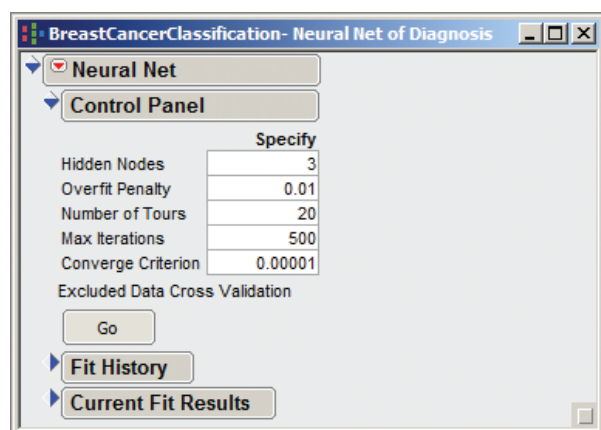


Figure 25. Neural Net Control Panel settings

Results for the current fit (JMP provides a **Fit History** panel that retains history for all fits) are shown in the **Current Fit Results** panel (Figure 26). In our sequence of 20 tours – recall that each has a random start – only one of the 20 tours **Converged at Best**, and this is the fit that JMP uses. (We have saved the script for this analysis in **Neural Net Report**. However, because of the random starts, when you run this script, your results will likely differ from those in Figure 26.) Note that 97 parameters (**Nparm**) have been estimated. These are listed in the **Parameter Estimates** panel, but they are not of intrinsic interest.

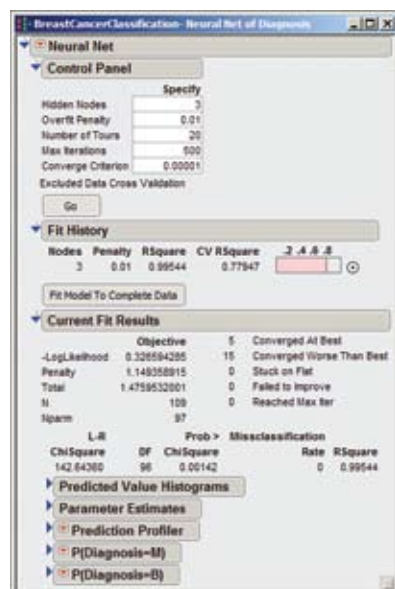


Figure 26. Results of neural net fit to diagnosis with 30 predictors

In order to be able to refer to the same model and same results, we have saved a script to the data table that reproduces the model that we explore below (the one that was selected in Figure 26). This script is called **Neural Net 1**. If you are following using JMP, you should run this script at this point.

After running **Neural Net 1**, under the red triangle, we choose the option to **Save Formulas**. This adds five new columns to the data table: the formulas for the three hidden nodes, called **H1 Formula**, **H2 Formula**, and **H3 Formula**; a formula called **SoftSum**, which calculates an intermediate result; and the final estimated probability, in our case called **Diagnosis[M] Formula**, which applies a logistic function to the estimated linear function of the hidden nodes.

Histograms for the probabilities predicted by the model, conditioned on whether the actual **Diagnosis** is malignant or benign, are provided in the Neural Net report, as shown in Figure 27. Note that, for the training data, the probabilities predicted by the model appear to classify masses into the two classes quite accurately. Also note that the **RSquare** for the model, fit to the training data, is 0.99547.

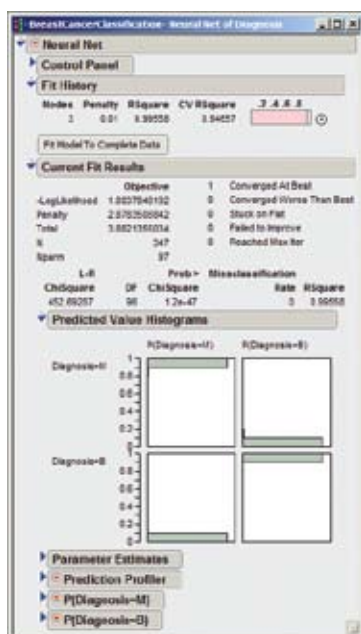


Figure 27. Histogram for neural net estimate of probability of malignant

To get some sense of what our 31-dimensional model looks like, the Neural Net report provides various options. Under the red triangle in the Neural Net report window, we can choose **Profiler**. This opens the **Prediction Profiler**, a portion of which is shown in Figure 28. The dotted red vertical lines represent settings for each of the variables. For a given variable, the traces in the **Prob(Diagnosis=M)** and **Prob(Diagnosis=B)** squares represent the cross sections of the fitted model in that variable's direction, at the settings of the other variables. When we change one variable's value, we can see the impact of this change on the surface for all other variables. As one changes values and scrolls through the plots, one sees that the surface appears fairly smooth, with some steep peaks, but no very jagged areas.

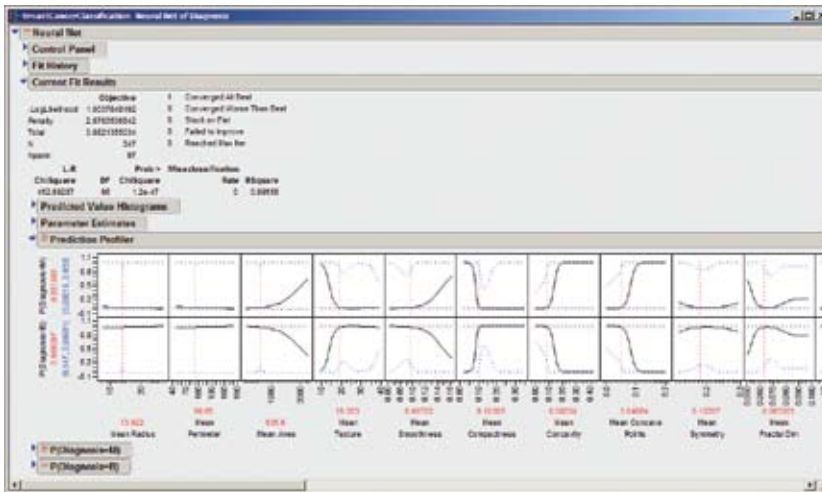


Figure 28. Part of prediction profiler for Neural Net Model 1

Another way to visualize the surface is using the **Surface Profiler**. This gives a three-dimensional view of the effect of predictor variables, taken two at a time, on $\text{Prob}(\text{Diagnosis}=\text{M})$ and $\text{Prob}(\text{Diagnosis}=\text{B})$. Figure 29 shows one of the $30 \times 29/2 = 435$ possible surface plots for $\text{Prob}(\text{Diagnosis}=\text{M})$.

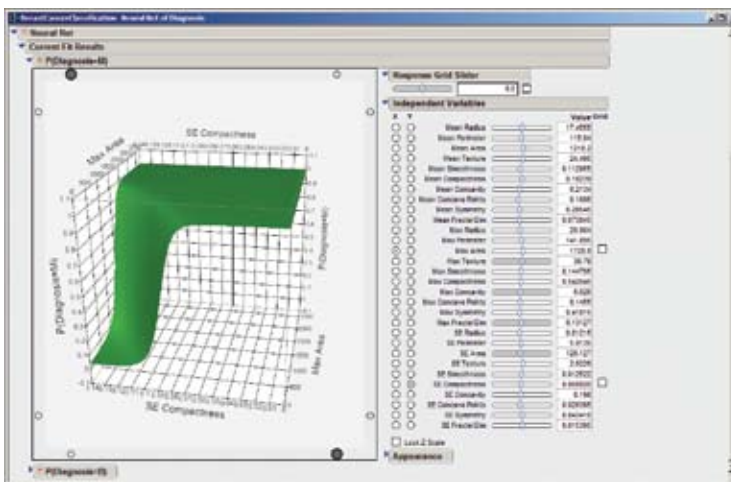


Figure 29. Surface plot of probability of malignancy as a function of Max Area and SE Compactness in Neural Net Model 1

To assess the performance of this model on the training set, we define a variable called **NN1 Prediction**. As usual, predicted probabilities equal to or greater than 0.50 will lead to a classification of M, while values less than 0.50 will result in a classification of B. The resulting **Fit Y by X** analysis, for the training set, is shown in Figure 30. There are no misclassifications. But keep in mind that neural nets have a tendency to overfit the training set.

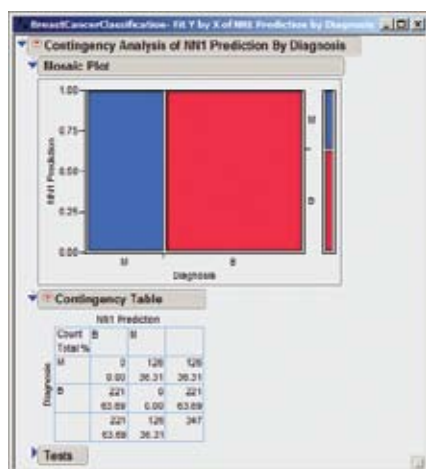


Figure 30. Mosaic Plot and Contingency Table for classification based on Neural Net Model 1

In the neural net platform, one can explore the effects of various variable combinations and specifications of hidden layer nodes. For example, if one uses the variables from our logistic stepwise selection and a single hidden node, then the resulting neural net model gives predicted probabilities that are similar to those obtained using the logistic fit. One can also utilize crossvalidation to determine a model that might generalize well to new data. This is the topic of the next section.

7.3 Second model – Neural Net 2

In this section, we will explore various neural net model architectures using K-fold crossvalidation. Once again, we go to **Analyze > Modeling > Neural Net**. We fill in the launch window as before, with **Diagnosis** as **Y** and all 30 predictors as our **Xs**. However, we make another selection; namely, we choose **K-Fold Crossvalidation** from the drop-down menu at the bottom of the launch window, as shown in Figure 31. (We note that we will be doing crossvalidation within the context of our 347 observation training set, in order to have consistency with our other three models. In practice, when using K-fold crossvalidation, we would not hold out a separate validation set.)

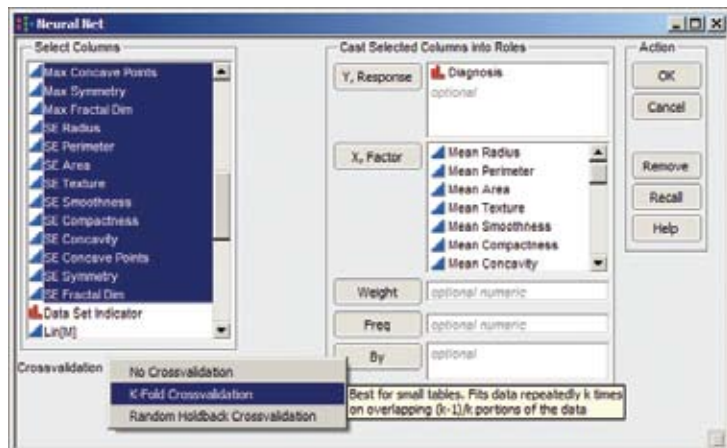


Figure 31. Launch window for neural net with K-fold crossvalidation

Clicking **OK** brings up the **Control Panel**. The default number of groups is five, and we consider that reasonable, given the size of our data set – we have 347 observations, and so each of the five training samples will contain at least 276 rows (there are about $347/5 = 69$ observations in each fold, and four folds in the training sample, so about $69 \times 4 = 276$ observations in each training set). When we click **Go**, the **Control Panel** changes back to give us the choices that we have seen earlier.

This time, we will explore the effect of different numbers of hidden nodes and various values of the overfit penalty on the fit. To do this, we select **Sequence of Fits** from the red triangle drop-down menu in the Neural Net report. We populate the resulting dialog box as shown in Figure 32. This will allow us to explore models with two, three and four hidden nodes, and overfit penalties that vary from 0.01 to 1.28, by multiples of two. This is a total of 24 models.

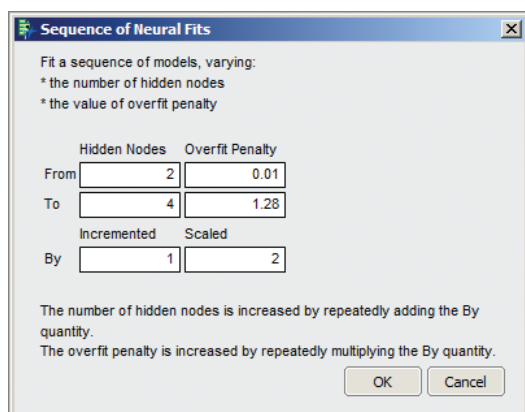


Figure 32. Sequence of Neural Fits dialog box

The models that JMP calculates in the sequence of fits are summarized in the **Fit History** panel. Again, keep in mind that, because of random starts, your results will not exactly match ours. Our results are shown in the **Fit History** panel in Figure 33. We note that there are three models with a **CV RSquare** of 0.85 or better. Two of these have four nodes (and so 129 parameters) and one has three nodes (97 parameters). Since simpler models are almost always more desirable, we choose the three-parameter model, with an overfit penalty of 0.64, as our current model. We click the radio button to the right of the model and the **Current Fit Results** panel updates to show summary values for that model (see Figure 33.)

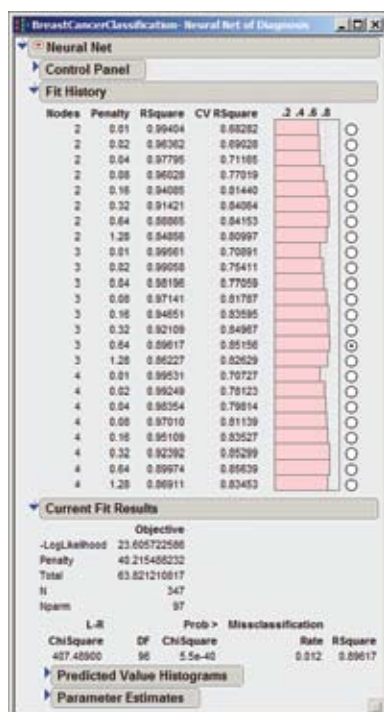


Figure 33. Fit History panel for 24 exploratory model fits

As before, because of the random starts, the results that we obtain likely differ from yours. However, the script **Neural Net 2** reproduces the final model that we chose in Figure 33.

If you are following along, please run **Neural Net 2**. Under the report's red triangle, we click on **Profiler** and **Surface Profiler** to get a sense of what this model looks like. When we do this, we see that this model seems considerably smoother than our **Neural Net 1** model. For comparison, we show some of the profiler traces in Figure 34.



Figure 34. Part of Prediction Profiler for Neural Net Model 2

As before, to save this model's formulas we choose **Save Formulas** from under the report's red triangle. This inserts five new columns into the data table: **H1 Formula 2**, **H2 Formula 2**, **H3 Formula 2**, **SoftSum 2** and **Diagnosis[M] Formula 2**. We define a new column called **NN2 Prediction** that gives the classifications for Neural Net 2. The **Fit Y by X** analysis in Figure 35 shows four misclassifications. Recall that there were no misclassifications for Neural Net 1 on the training data (Figure 30). Did Neural Net Model 1 overfit the data? We will keep that in mind during the model validation step.

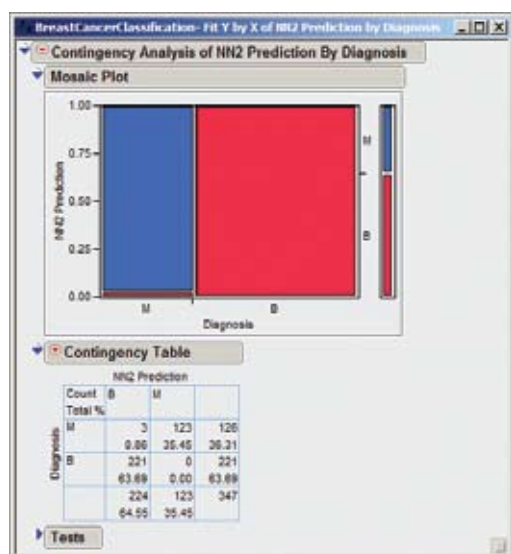


Figure 35. Mosaic Plot and Contingency Table for classification based on Neural Net Model 2

8. Comparison of classification models

At this point, we will select a single model by comparing the performance of all four models on our validation set. This kind of activity is sometimes referred to as “competing models.” Our single performance criterion will be the overall misclassification rate. But note that, in a real setting, we might want to weight the two types of misclassification errors differently, since overlooking a malignant tumor seems more serious than incorrectly classifying a benign tumor. Although JMP provides mechanisms for accomplishing this (for example, one could define a loss function and fit models using the nonlinear platform), we will not delve into this here.

We can easily switch our analyses to our validation set – simply locate the row state variable **Validation Set** in the **Columns** panel of the data table, left-click on the red star to its left, and select **Copy to Row States**. All but the 109 observations in our validation set are now excluded and hidden. We go to **Analyze > Fit Y by X**, and we enter as **Y, Response** the variables **Most Likely Diagnosis** (the logistic classification), **Partition Prediction**, **NN1 Prediction** and **NN2 Prediction**. As **X, Factor**, we enter **Diagnosis** (the actual diagnosis for that tumor). The script is saved as **Performance Comparison**.

The report is shown in Figure 36. The logistic and neural net models outperform the partition model. Both neural net models slightly outperform the logistic model. Neural Net 2 slightly outperforms Neural Net 1. To choose between these two neural net models, we refer back to the sequence of fits analysis shown in Figure 33. We note that Neural Net 1, which has three nodes and an overfit penalty of 0.01, is represented, in the sequence of fits, by a neural net model that has a crossvalidation RSquare of 0.70891. For this data, larger overfit penalties seem to result in better crossvalidation results. For this reason, we prefer Neural Net 2 over Neural Net 1. On this basis, we choose Neural Net Model 2 as our classifier.




Figure 36. Comparison of four models on validation data

Now, to obtain a sense of how this model will perform on new, independent data, we apply it to our test set. We apply row states from the row state variable **Test Set** to the data table in order to exclude all but the test set observations. We rerun the script **Performance Comparison** to see the performance of Neural Net Model 2 on the test data. Although the report (Figure 37) shows results for all four models, it is Neural Net Model 2 that is of primary interest at this point, as it is our chosen model. As hoped for, our chosen model outperforms the others on the test data, netting us a misclassification rate of 3/113.




Figure 37. Performance of four models on test data

9. Conclusion

The goal of this paper was to illustrate some of the features of JMP that support classification and data mining. We began by illustrating various visualization techniques that provide an understanding of the data and relationships among the variables. We partitioned our data into a training set, a validation set and a test set. We then fit four models using the training data: a logistic model, a partition model and two neural net models. The best classification, based on performance on the validation set, was obtained using a neural net model whose structure was chosen using K-fold crossvalidation.

We note that the logistic and neural net models had similar performance. Partition models tend not to perform as well as nonlinear (or linear) regression techniques when the predictors are continuous. They can be very useful, though, when there are categorical predictors, and especially when these have many levels. And, unlike neural net models and even logistic models, partition models are very intuitive and interpretable. In our situation, where classification was the primary goal, the interpretability of the model was less important than its ability to classify accurately.

We also wish to underscore the importance of guarding against overfitting, which, in the case of neural net models, often results in claims of exaggerated model performance. The application of K-fold crossvalidation helped us arrive at a neural model that was simple and that generalized well to our test set. Also, in the case of neural nets, where overfitting is so easy, it is important to assess the model's performance on an independent data set. Without this step, claims about model performance risk being exaggerated.

References

Mangasarian, OL, Street, WN, and Wolberg, WH, "Breast Cancer Diagnosis and Prognosis via Linear Programming," Mathematical Programming Technical Report 94-10, Dec. 19, 1994, pp. 1-9.

Street, WN, Wolberg, WH, and Mangasarian, OL, "Nuclear Feature Extraction for Breast Tumor Diagnosis," 1993 International Symposium on Electronic Imaging: Science and Technology, Vol. 1905, 1993, pp. 861-870.

Bishop, CM, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.

Ripley, BD, *Pattern Recognition and Neural Networks*, Cambridge University Press, 1996.

McCulloch, WS and Pitts, WH, "A Logical Calculus of the Ideas Immanent in Nervous Activity," Bull. Math. Biophys. 5, 115-133, 1943.

Acknowledgments

Mia Stephens and Phil Ramsey, both partners in the North Haven Group, made significant contributions to this paper. Christopher Gotwalt at SAS was extremely generous with details on JMP software's implementation of neural nets. Ian Cox, also at SAS, reviewed this paper with care and made many useful suggestions. Many thanks are due to each of these individuals.

Portions of this work are taken from the SAS Press publication: *Visual Six Sigma: Making Data Analysis Lean*.

Recommended Reading

Berry, MJA and Linoff, GS, *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management*, John Wiley and Sons, 1997.

Hastie, T, Tibshirani, R, and Friedman, J, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, 2001.

Haykin, S, *Neural Networks: A Comprehensive Foundation*, 2nd Edition, Prentice Hall, 1998.

Sall, JP and Maahs-Fladung, C, "Trees, Neural Nets, PLS, I-Optimal Designs and Other New JMP® Version 5 Novelties", SUGI 27, <http://www2.sas.com/proceedings/sugi27/p268-27.pdf>.

Contact Information

Your comments and questions are valued and encouraged. Contact the author at:

Name..... Marie Gaudard

Enterprise..... North Haven Group

Address..... PO Box 296

City, State ZIP Hernando, FL 34442

Work Phone: 352-560-0312

E-mail: mgaudard@northhavengroup.com

Web: www.northhavengroup.com



SAS Institute Inc. World Headquarters +1 919 677 8000 To contact your local JMP office, please visit www.jmp.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. Copyright © 2008, SAS Institute Inc. All rights reserved. 103723_520006.1108