

Using JMP to Compare Models from Various Environments

Lucas Beverlin

October 15, 2020

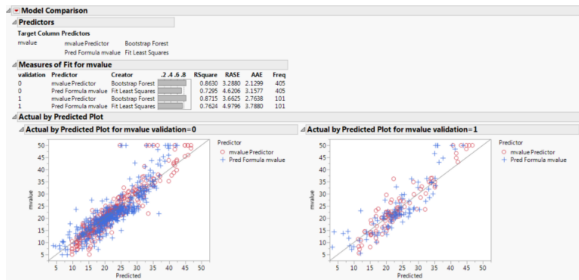
JMP Can Do Many Things

JMP 15 Pro has several platforms that can build a model:

- Fit Model
- Neural
- Partition
- Nonlinear
- And more!

And Compare the Models

The Model Comparison platform can be used to compare fits from various JMP platforms, to help the user determine which model is best.



(Figure 10.1 from the Predictive and Specialized Modeling pdf within the JMP 15 Pro help.)

But It Can't Do Everything

There are some models that JMP cannot fit, and some limitations to those that it can:

- A neural network with more than two hidden layers, e.g. autoencoders, convolutional neural networks, etc., and/or activation functions other than the usual 3 JMP Pro offers
- Projection pursuit regression
- Multivariate adaptive regression splines (MARS)

Software That Can Fit Other Models

- R
- Python
- MATLAB

Of course, there are many more, but for this talk we'll focus on these three. The ideas here can easily extend to basically any software!

The Power of Model Comparison

We can use the Model Comparison platform to compare model fits from other software too!

What you need:

- The model predictions from whatever you used to fit the model.
- If you split the data into training/validation/test sets, something that delineates those sets.

Calling R from JMP

JMP's scripting language (JSL) can create its own R session and run code from it.

```
dt = Open("C:\Program Files\SAS\JMPPRO\14\Samples\Data\Boston Housing.jmp");
Set Environment Variable( "R_HOME", "C:\Program Files\R\R-3.6.3" );//Handy if you have multiple versions of R!
R Init();
R Send(dt);
R Submit( "
lm1<-lm(y~.,data=dt)
preds<-predict(lm1,dt)
" );
R Get(preds)
```

<https://www.r-project.org>

Notes: I am using R 3.6.3 for this demonstration. JSL can call any version of R $i = 2.9.1$

Calling Python from JMP

JSL can create its own Python session and run code from it.

```
dt = Open("C:\Program Files\SAS\JMPPRO\14\Samples\Data\Boston Housing.jmp");  
  
Python Init();  
Python Send(dt);  
Python Submit("  
import numpy as np  
  
X = dt.iloc[:, :-1].values  
y = dt.iloc[:, 1].values  
  
from sklearn.linear_model import LinearRegression  
regressor = LinearRegression()  
regressor.fit(X, y)  
  
y_pred = regressor.predict(X)  
  
");  
Python Get(y_pred)
```

<https://python.org>

Notes: While JSL can call the newest version of Python, which at the time of recording was 3.8.5, it cannot call versions of Python greater than 3.6.x if you use the Anaconda installer. Thus, I am using Python 3.6.5 for this demonstration. JMP recommends

Calling MATLAB from JMP

JSL can create its own MATLAB session and run code from it.

```
dt = Open("C:\Program Files\SAS\JMPPRO\14\Samples\Data\Boston Housing.jmp");  
  
Set Environment Variable( "MATLABROOT", "C:\MATLAB\R2019b");  
  
dtm = dt<<Get as Matrix;  
MATLAB Init();  
MATLAB Send(dtm);  
MATLAB Submit("  
  
    X=dtm(:,1:13)  
    y=dtm(:,14)  
    lsfit = fitlm(X,y,'Categorical',[4])  
    preds = predict(lsfit,X)  
  
");  
  
MATLAB Get(preds)
```

<https://www.mathworks.com/products/matlab.html>

Notes: I am using Matlab R2019b for this demonstration. JSL can call any version of MATLAB \geq R2012a (7.14.0).

Using JSL to Fit Models and Run Model Comparison

So how do we tie this all together?

- ① Fit each model, sending each software the data, and which set (training/validation/test) each observation resides.
- ② Output the model fits, and add them to the JMP data table. (Be sure to name them so that you can tell one from the other!)
- ③ Depending on the model, you may want some model diagnostics. Those can be outputted as well.
- ④ Use the Model Comparison platform to compare all the fits, and it will present tables for output showing the best in each set (training/validation/test).

Model Fitting

For the first step, we can do one of two things:

- 1 Tell JMP via JSL to call your software of choice, send it the data, and the code to fit it.
- 2 Create a dataset with a validation column, and schedule it to run on your software of choice, outputting what you need (predictions, diagnostic plots, etc.).

Warnings:

- Make sure all software has what it needs to fit the model!
- Make sure all software has finished running before proceeding!
- Make sure the output format is something JMP can read!
- Make sure the predictions from each model correspond to the correct observations from the original data set!

What I Used to Fit Models

In R I used the following additional packages to fit a neural network:

- Keras
- Tensorflow

In Python, I used the following libraries:

- Keras
- Tensorflow
- numpy - to do a few calculations
- pandas - to manipulate the data a bit
- matplotlib - to create plots

In MATLAB, I used the deep learning toolbox.

Adding Predictions to the JMP Data Table

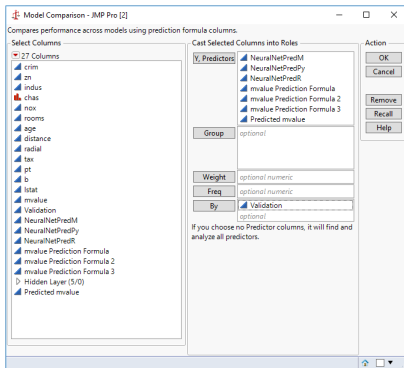
- 1 If you used JMP to call the software, you can use JSL code to retrieve the predictions and add them to the data table.
- 2 If the software ran on its own, and saved the output, you can ask JMP to read the output, and add it as a new column. We can also read the diagnostic plots, assuming they were saved as graphics files.

To help with identification later, I recommend adding the following property for each prediction column you add:

```
dt << Set Property("Predicting",{:mvalue, Creator("R")})
```

Here this will call the model's creator as R in the output. Of course, choose whatever string you wish to identify each model, particularly if multiple models come from the same software.

Model Comparison



Note that you can also use the Validation column here as a Group variable instead of a By variable.

Model Comparison Output

Training Diagnostics									
Predictors									
Measures of Fit for mvalue									
Predictor	Creator	.2	.4	.6	.8	RSquare	RASE	AAE	Freq
NeuralNetPredM	MATLAB					0.7822	4.5525	3.0060	304
NeuralNetPredPy	Python					0.7623	4.7561	3.8669	304
NeuralNetPredR	R					0.6111	6.0837	4.2576	304
mvalue Prediction Formula	Fit Generalized Standard Least Squares					0.7464	4.9129	3.4574	304
mvalue Prediction Formula 2	Fit Generalized Lasso					0.7372	5.0008	3.4305	304
mvalue Prediction Formula 3	Fit Generalized Ridge					0.7295	5.0743	3.4247	304
Predicted mvalue	Neural					0.7781	4.5952	3.3754	304

Validation Diagnostics									
Predictors									
Measures of Fit for mvalue									
Predictor	Creator	.2	.4	.6	.8	RSquare	RASE	AAE	Freq
NeuralNetPredM	MATLAB					0.7890	3.9142	2.8937	101
NeuralNetPredPy	Python					0.7171	4.5324	3.3725	101
NeuralNetPredR	R					0.6044	5.3600	3.5845	101
mvalue Prediction Formula	Fit Generalized Standard Least Squares					0.7077	4.6070	3.4992	101
mvalue Prediction Formula 2	Fit Generalized Lasso					0.7198	4.5105	3.3940	101
mvalue Prediction Formula 3	Fit Generalized Ridge					0.7339	4.3954	3.2904	101
Predicted mvalue	Neural					0.7853	3.9483	3.0428	101

Test Diagnostics									
Predictors									
Measures of Fit for mvalue									
Predictor	Creator	.2	.4	.6	.8	RSquare	RASE	AAE	Freq
NeuralNetPredM	MATLAB					0.6542	4.6219	3.1106	101
NeuralNetPredPy	Python					0.8027	3.4916	2.8458	101
NeuralNetPredR	R					0.6171	4.8637	3.5895	101
mvalue Prediction Formula	Fit Generalized Standard Least Squares					0.6757	4.4761	3.3397	101
mvalue Prediction Formula 2	Fit Generalized Lasso					0.6898	4.3778	3.2233	101
mvalue Prediction Formula 3	Fit Generalized Ridge					0.7067	4.2565	3.1275	101
Predicted mvalue	Neural					0.6993	4.3100	3.2619	101

Here, had you used Validation as a Group variable, this would be one table instead of three.

Example: Boston Housing


We will use the classic Boston Housing dataset. (Harrison & Rubinfeld 1978)

- mvalue - Median value of a house in thousands
- crim - Per capita crime rate by town
- zn - Proportion of residential land zoned for lots over 25,000 ft²
- indus - Proportion of non-retail business acres per town
- chas - 1 if the tract borders the Charles River, and 0 otherwise
- nox - NO_x concentration (parts per 10 million)
- rooms - Average number of rooms per dwelling
- age - Proportion of owner-occupied units built before 1940
- distance - Weighted distance to 5 Boston employment centers
- radial - Index of accessibility to radial highways
- tax - Full-value property tax rate per \$10,000
- pt - Pupil/teacher ratio by town
- $b = 1000(B - 0.63)^2$, where B = proportion of African Americans by town
- lstat - % lower status of the population

Models to Consider

With the Model Comparison platform, we can compare all types of models to one dataset. In this presentation, we'll consider the following:

Linear Regression



Ridge Regression

Ridge regression is linear regression with the constraint that $\sum_{k=1}^m \beta_k^2 \leq c$, or, equivalently,

$$\hat{\beta} = \min_{\beta} \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{k=1}^m \beta_k x_{ki} \right)^2 + \lambda \sum_{k=1}^m \beta_k^2.$$

This method shrinks the coefficients toward 0.

The Lasso


The lasso is linear regression with the constraint that $\sum_{k=1}^m |\beta_k| \leq c$, or, equivalently,

$$\hat{\beta} = \min_{\beta} \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{k=1}^m \beta_k x_{ki} \right)^2 + \lambda \sum_{k=1}^m |\beta_k|.$$


This method can eliminate inputs whose contribution to the response is weak.

The Lasso and Ridge Regression

Lasso

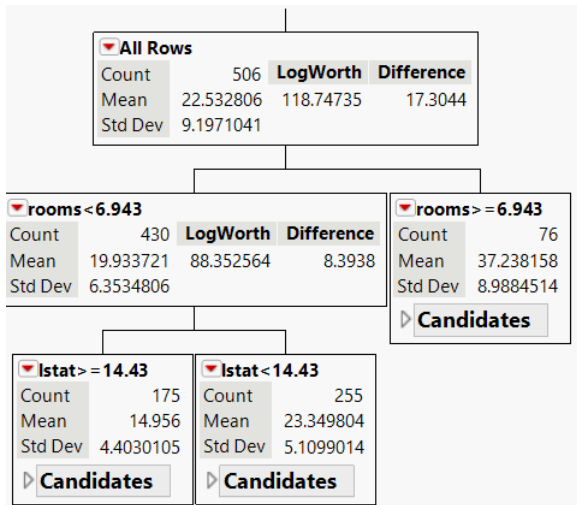


Ridge Regression




Thanks James, Witten, Hastie, and Tibshirani!

Regression Trees




Regression Trees

Regression trees are models where the input space is partitioned into areas, and the prediction is the average response of the observations in that area.



Thanks James, Witten, Tibshirani, and Friedman!

Neural Networks



Neural Networks

Neural networks are nonlinear models whose representation can be drawn as a series a neurons, much like how we believe the brain works.

- Each input can be passed to nodes in a hidden layer.
- At the hidden layer, the inputs are pushed through an activation function, and an output is calculated.
- Each output can be passed to a node in another hidden layer, or be an output of the network.

Notes on the Neural platform within JMP Pro:

- There are only three options for activation functions: Linear, Hyperbolic Tangent, and Gaussian Radial Basis.
- JMP Pro only allows for 2 hidden layers. Deep learning neural networks, such as autoencoders, often have dozens.

Projection Pursuit Regression

Projection pursuit regression, originally proposed by Friedman and Stuetzle (1981), is a model whose predictions take the form:

$$\mathbf{y} = \sum_{i=1}^k \beta_i f_i(\alpha' \mathbf{X}).$$

One can see that this is somewhat analogous to a neural network with one hidden layer with k nodes, each with activation function f_i . However, here, f_i is also estimated, as it is usually a smoother or a spline fit. (The f_i 's are called ridge functions.) Of course, by altering the α and β_i , the optimal smooth will change, so fitting is usually done stagewise.

Model Comparison in Action - JMP Journal Output

Do note that I had difficulty setting the seeds for the neural networks fit in Keras for R and Python, so some of your results will differ slightly from mine.

ModelComparisonOutput - JMP Pro

Training Diagnostics

Predictors

Measures of Fit for mvalue

Predictor	Creator	R Square	RMSE	AAE	Freq
NeuralNetPsdM	MATLAB	0.7822	4.5525	3.0090	304
NeuralNetPsdPy	Python	0.7823	4.5514	2.9812	304
NeuralNetPsdR	R	0.6334	5.9094	3.8279	304
mvalue Prediction Formula	Fit Generalized Standard Least Squares	0.7464	4.9129	3.4274	304
mvalue Prediction Formula 2	Fit Generalized Lasso	0.7372	5.0008	3.4305	304
mvalue Prediction Formula 3	Fit Generalized Ridge	0.7395	5.0743	3.4247	304
Predicted mvalue	Neural	0.7766	4.6705	3.3819	304
PPFPsdR	R	0.8906	3.2271	2.2389	304

Validation Diagnostics

Predictors

Measures of Fit for mvalue

Predictor	Creator	R Square	RMSE	AAE	Freq
NeuralNetPsdM	MATLAB	0.7890	3.9142	2.8937	101
NeuralNetPsdPy	Python	0.8025	3.7583	2.8355	101
NeuralNetPsdR	R	0.7373	4.3873	3.5232	101
mvalue Prediction Formula	Fit Generalized Standard Least Squares	0.7077	4.6070	3.4992	101
mvalue Prediction Formula 2	Fit Generalized Lasso	0.7196	4.5105	3.3940	101
mvalue Prediction Formula 3	Fit Generalized Ridge	0.7339	4.3954	3.2904	101
Predicted mvalue	Neural	0.7681	3.9329	3.0332	101
PPFPsdR	R	0.8480	3.3444	2.5204	101

Test Diagnostics

Predictors

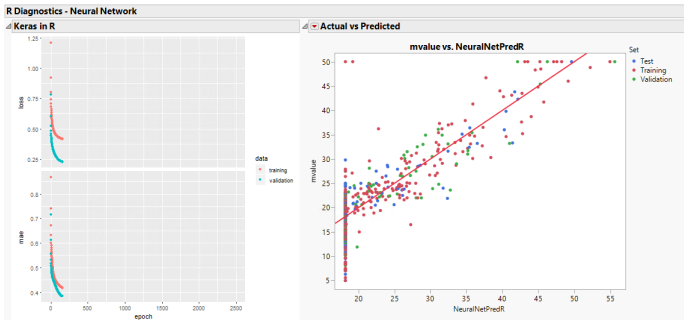
Measures of Fit for mvalue

Predictor	Creator	R Square	RMSE	AAE	Freq
NeuralNetPsdM	MATLAB	0.6542	4.6219	3.1106	101
NeuralNetPsdPy	Python	0.6746	4.4823	3.1188	101
NeuralNetPsdR	R	0.4635	4.5391	3.0515	101
mvalue Prediction Formula	Fit Generalized Standard Least Squares	0.6757	4.4761	3.3397	101
mvalue Prediction Formula 2	Fit Generalized Lasso	0.6896	4.3779	3.2239	101
mvalue Prediction Formula 3	Fit Generalized Ridge	0.7067	4.2365	3.1275	101
Predicted mvalue	Neural	0.7002	4.2385	3.1659	101
PPFPsdR	R	0.8324	4.7852	2.9788	101


Diagnostics

- R Diagnostics - Neural Network
- Python Diagnostics - Neural Network
- MATLAB Diagnostics - Neural Network
- Linear Regression Diagnostics
- Lasso Diagnostics
- Ridge Diagnostics
- R Diagnostics - Projection Pursuit

Model Comparison in Action - R Neural Network Diagnostic Plot




Model Comparison in Action - R Projection Pursuit Ridge Function Plots




Model Comparison in Action - Python Diagnostic Plot

Python Diagnostics - Neural Network


Keras in Python




Actual vs Predicted





Model Comparison in Action - MATLAB Diagnostic Plot





Model Comparison in Action - Linear Regression Diagnostic Plots and Output



Model Comparison in Action - Lasso Diagnostic Plots and Output



Model Comparison in Action - Ridge Regression Diagnostic Plots and Output



Which Model Fit Best?

- Here we appear to see that the projection pursuit model fit the training data the best, as well as the validation data. However, it was worst on the test data!
- It appears that JMP's neural network fit the test data best.
- Here, I would recommend JMP's neural network since the fit is fairly consistent across all three sets of data.

Conclusions

- The Model Comparison platform is a powerful tool for comparing the predictive ability of multiple models.
- This flexibility can be extended to include models that were not fit with JMP!
- This can allow for several complex models to be fit simultaneously, and once the results are in, they can all be compared.

References

Friedman, J. and Stuetzle, W. "Projection Pursuit Regression,"
Journal of the American Statistical Association, Vol. 76, 817-823,
1981.

Harrison, D. and Rubinfeld, D.L. "Hedonic Prices and the Demand
for Clear Air," J Environ. Economics & Management, Vol. 5,
81-102, 1978.

Hastie, T., Tibshirani R. and Friedman J. The Elements of
Statistical Learning, Springer, New York, 12th printing, 2017.

James, G., Witten, D., Hastie, T. and Tibshirani, R. An
Introduction to Statistical Learning with Applications in R.
Springer, New York, 8th printing, 2017.