University of Antwerp

ALPA IN WEKA

A quick introduction

Enric Junqué de Fortuny
David Martens
Faculty of Applied Economics,
University of Antwerp, Belgium

ALPA IN WEKA: A QUICK INTRODUCTION

What is the ALPA toolbox?

The ALPA toolbox provides algorithms to perform rule extraction on existing black-box models, this can serve two purposes:

1. Gaining insight into the underlying black-box for validation
2. Generating superior rule sets based on insights from a black-box

Installing ALPA

Follow these steps to install ALPA into an existing WEKA installation

1. Start WEKA version greater than 3.7 (GUI Chooser)
2. Menu [Tools] -> Menu Item [Package manager]
3. In the [Unofficial] panel on the right, select [File/URL]
4. Browse to the location of the ALPA ZIP file and select it.
5. Press OK
6. Restart WEKA (ALPA is now available under classifiers.rules)

A quick overview of this tutorial

Using ALPA is always a two step process: first you must identify a well performing black-box model set-up. Next you choose a desired white-box algorithm to explain your black-box and then run ALPA on both. This visual guide is split up in the following steps:

1. Load the data
2. Select and test the black-box
3. Select and test a white-box
4. Combine both with ALPA-R
5. Improving the comprehensibility

Note that basic working knowledge of WEKA is assumed, for more information see any of the many available WEKA tutorials or the official guide:

http://www.cs.waikato.ac.nz/ml/weka/book.html

The latest version of this document is always available at:

http://www.applieddatamining.com/files/ALPA/tutorial.zip

LOAD THE DATA

For this tutorial we will be using the Wisconsin Housing dataset, publicly available on the UCI machine learning repository and included with this tutorial. Locate the file *Wisconsin.arff* on your disk, it should have been included in the ZIP file containing this tutorial and can also be download from our website at:

http://www.applieddatamining.com/files/ALPA/tutorial.zip

Double click the file to open it in WEKA.

SELECT AND TEST THE BLACK-BOX

In the explorer, go to Classify->Choose and pick the SMOReg regressor (this is essentially a non-linear SVM by default). Non-linear SVMs are known to perform very well on a multitude of datasets, but the resulting models are often incomprehensible.

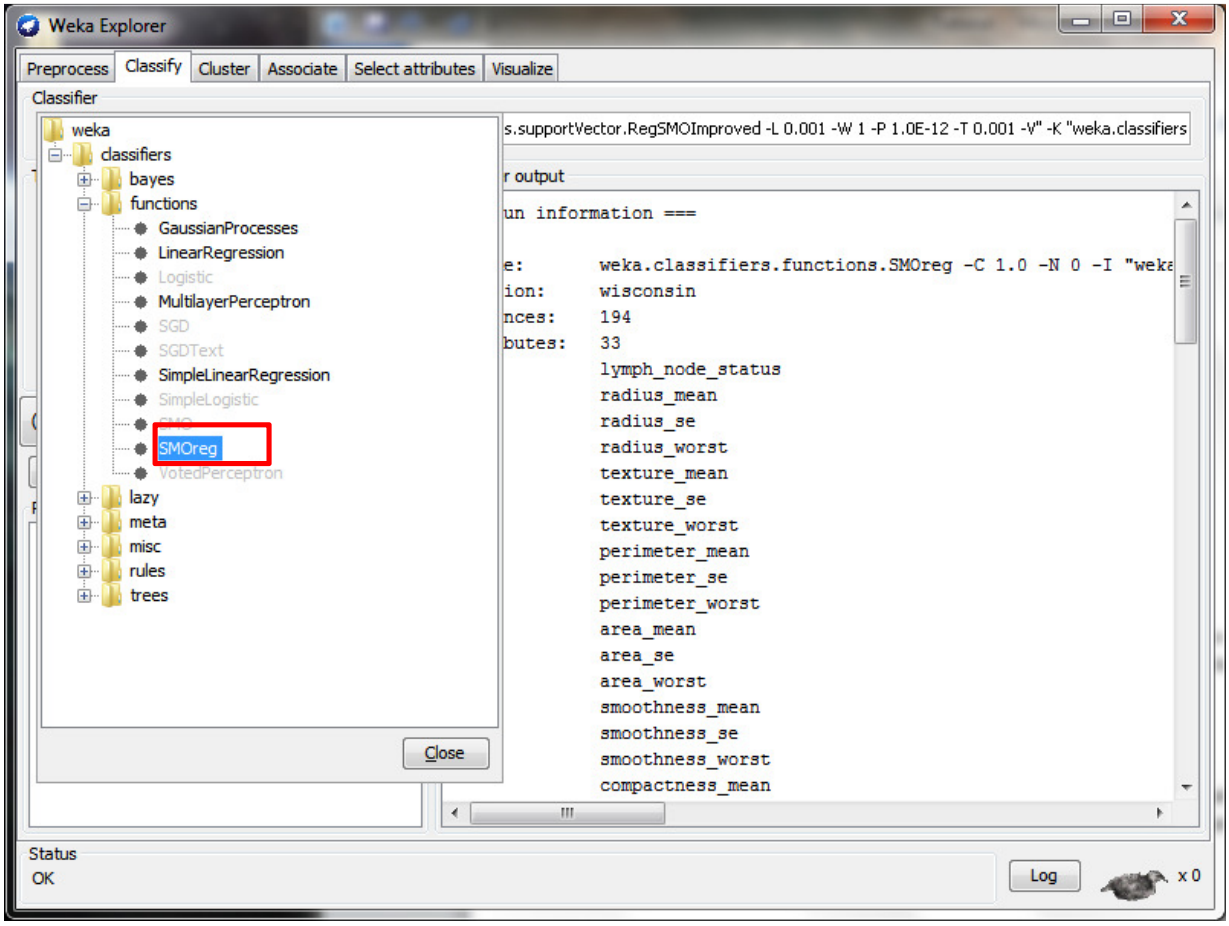

For the purpose of this document we will not tune the parameters but use the default setting with an RBF kernel as displayed below:

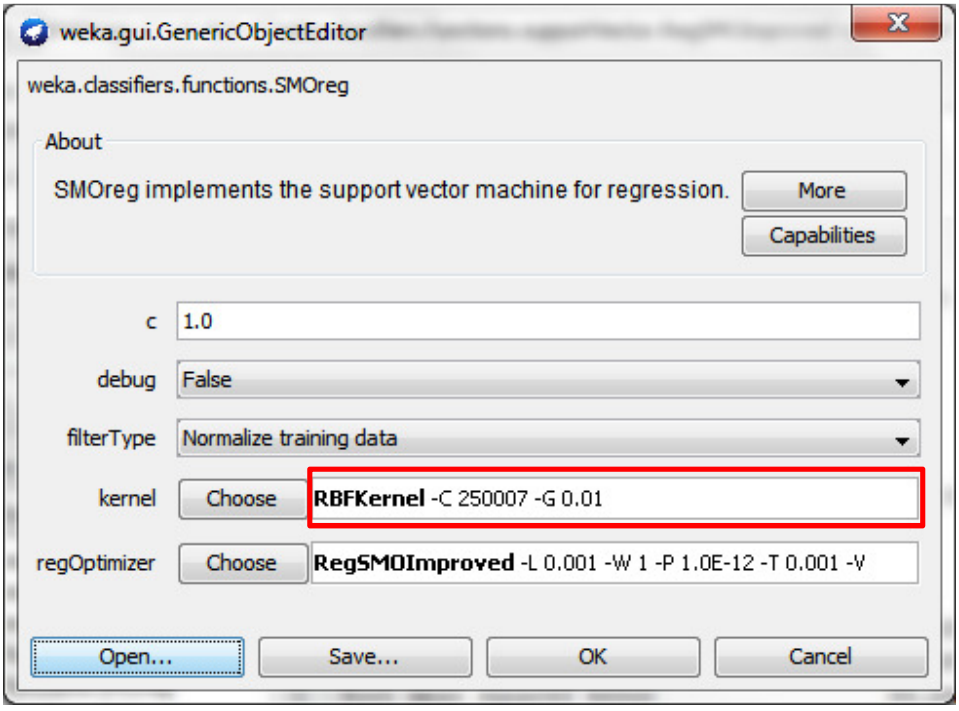

Select a test-split of 66% and start the experiment by clicking the start button. This should reveal an RMSE of about 31.74.

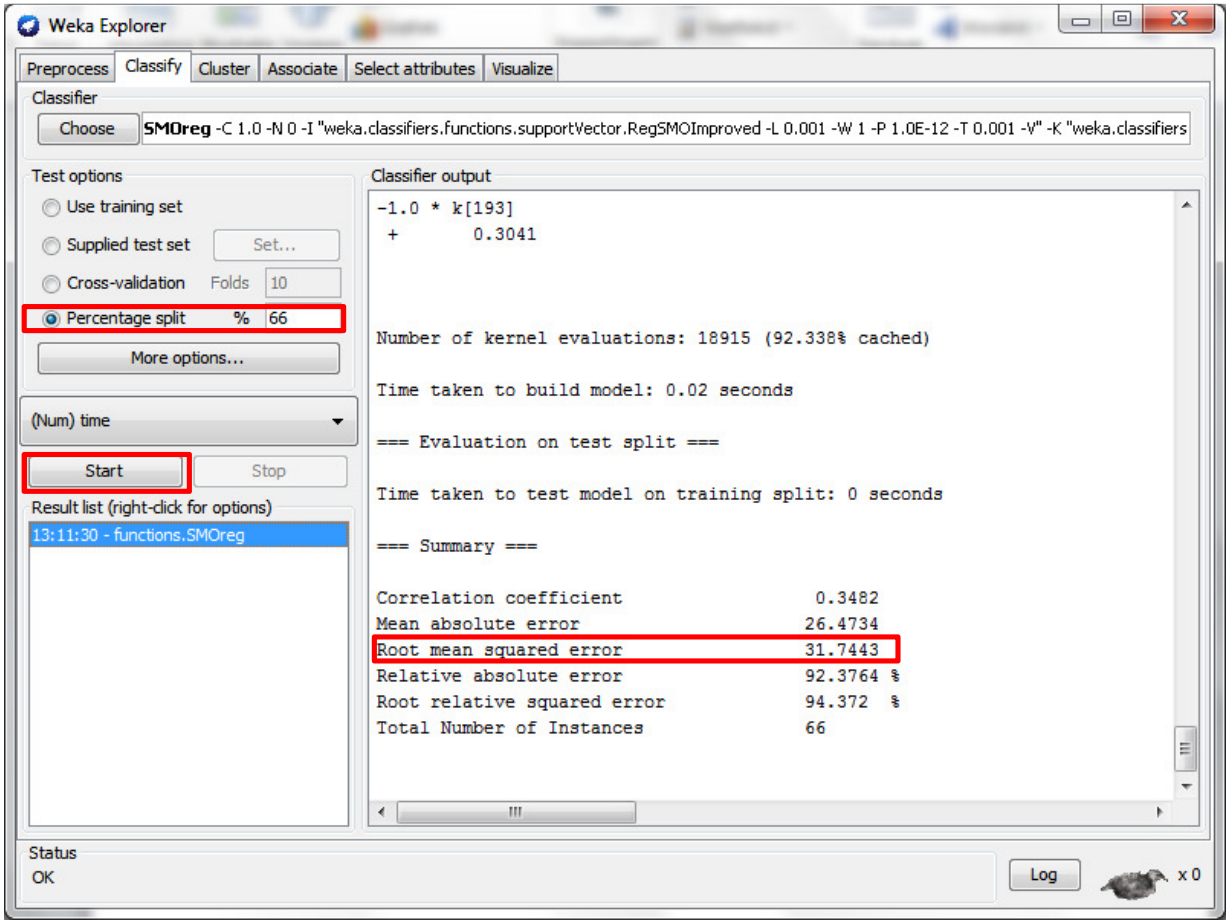

SELECT AND TEST THE WHITE-BOX

Let's check how our white-box model performs on this dataset, select the REPTree with a default configuration and repeat the experiment:

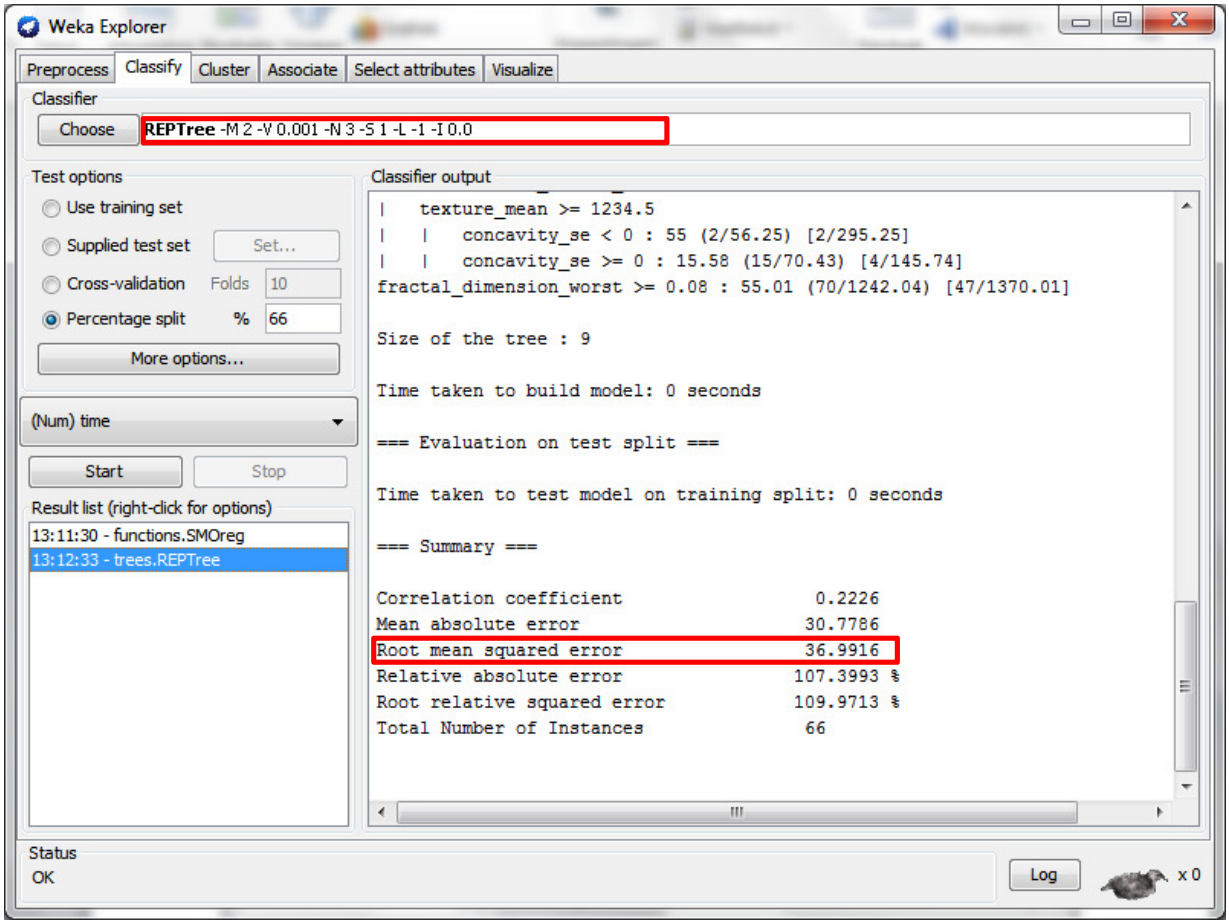

This RMSE of 36.99 is clearly worse than our black-box model's performance. Let's try using ALPA!

COMBINE BOTH IN ALPA

First select ALPA-R from the menu and select the same set-up as before for both the black-and the white-box. Make sure to select the right kernel for SMOReg like before.

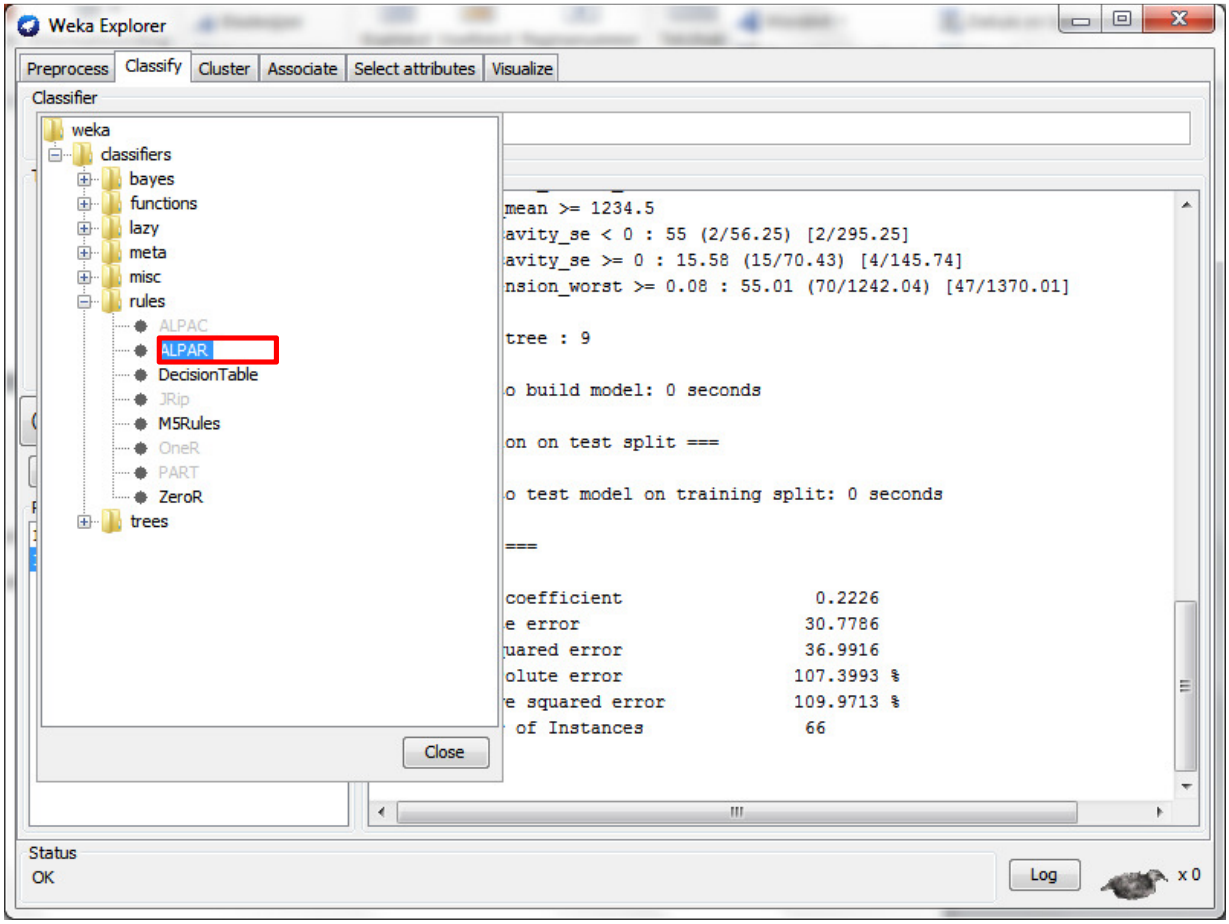

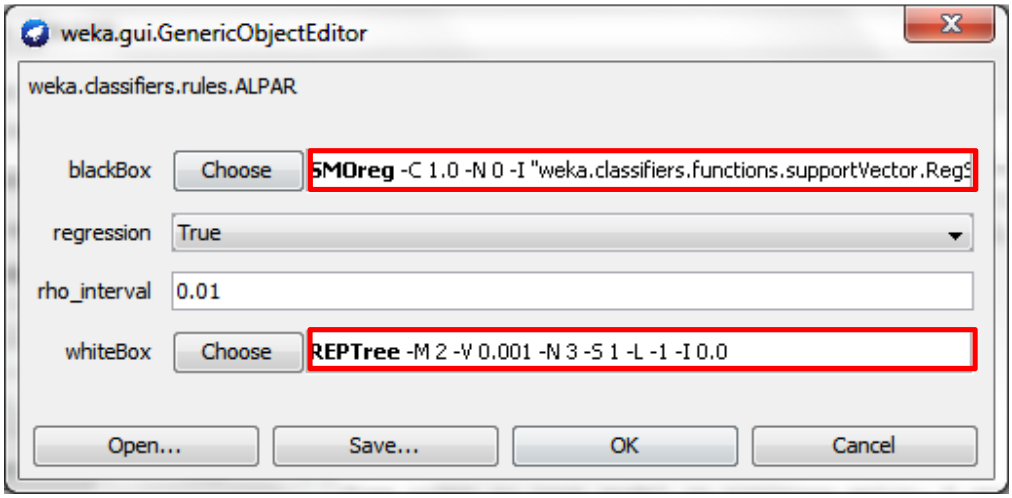

Click start, to start the training/testing procedure.

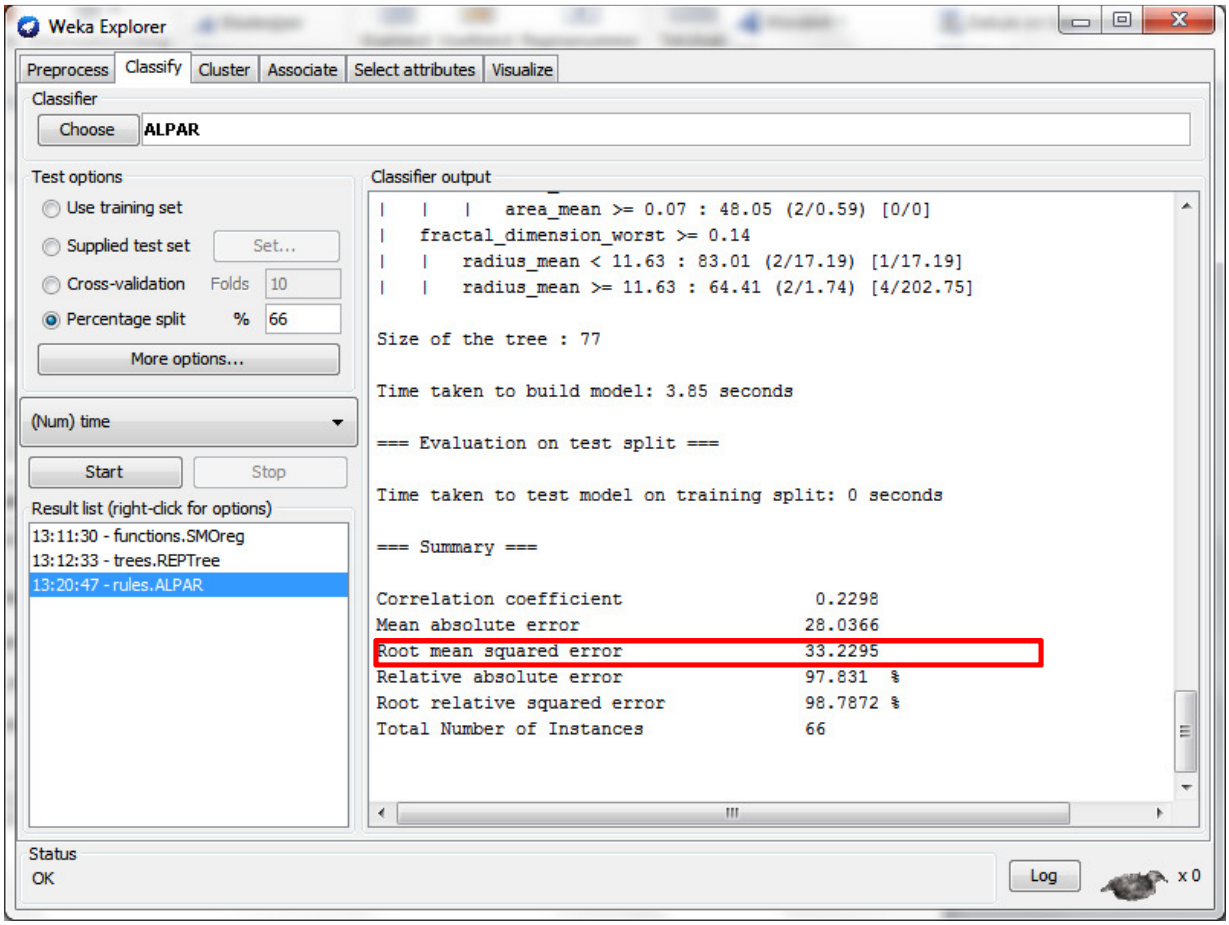

As is typical, the resulting model is a REPtree with *better performance* than the original REPTree. Note however, that this resulting tree is *quite large* and we might want to prune it down a little bit.

IMPROVING THE COMPREHENSIBILITY

We can make REPTrees more comprehensible by limiting the maximum depth of the resulting tree. Select a tree depth of three for this purpose in the ALPA settings → whiteBox (REPTree) → maxDepth and repeat the experiment.

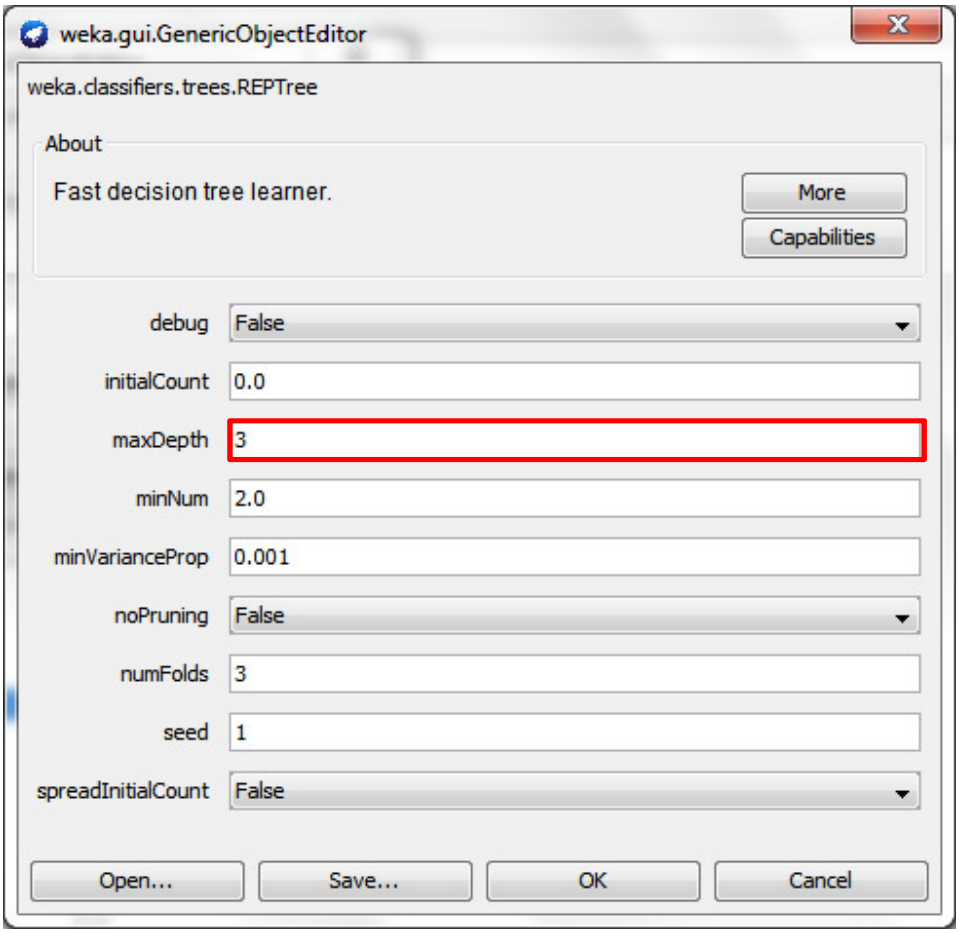

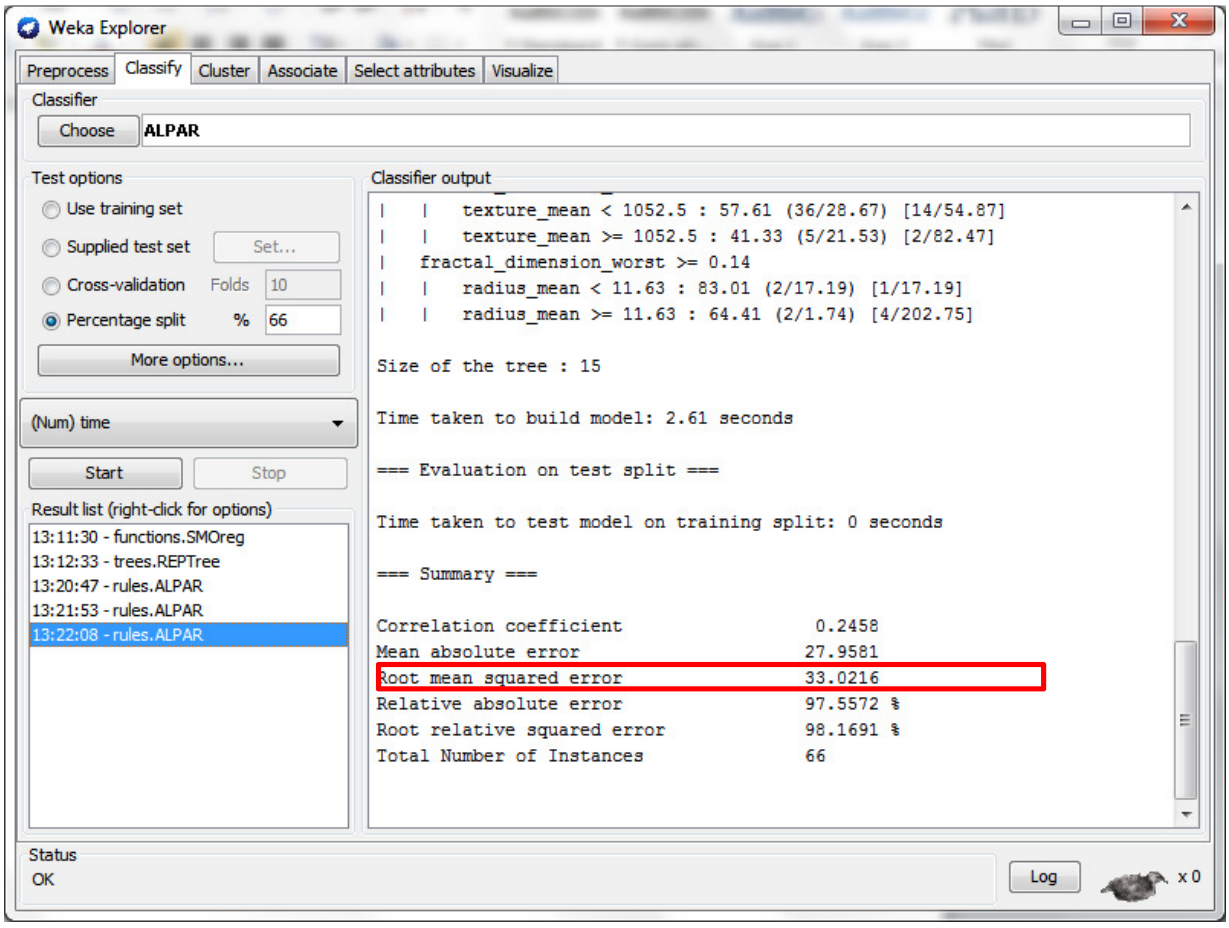

The resulting model still performs well but is much more comprehensible. You can visualize it by right clicking on rules.ALPAR and selecting “visualize tree”.

THE FINAL COMPREHENSIBLE MODEL

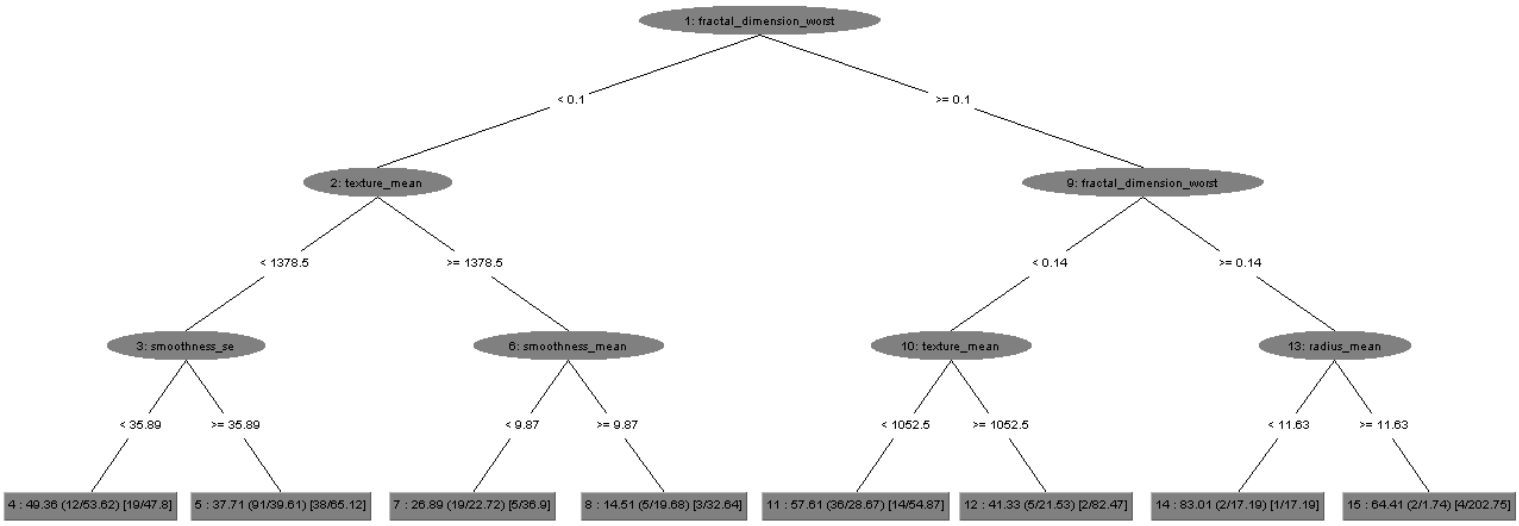

This model is comprehensible and has a much better performance than the original white-box. Furthermore, the output tells us that the validation fidelity is equal to 0.151, showing that it is a good explanation of the black-box model as well.

MORE INFORMATION

For more information about how to interpret the fidelity and RMSE values and the inner workings of ALPA, please refer to one of the following publications:

E Junqué de Fortuny, D Martens. *Active Learning-based Rule Extraction for Regression,* IEEE International Conference on Data Mining Workshop 2012

E Junqué de Fortuny, D Martens. *Active Learning-based Rule Extraction.* (Still under Review, available soon)

For bug reports, comments, suggestions or other inquiries, please contact me at:

Enric Junqué de Fortuny
Applied Data Mining Research Group,
Department of Engineering Management,
Faculty of Applied Economics,
Prinsstraat 13,
2000 Antwerp,
Belgium

enric.junquedefortuny@uantwerpen.be
http://www.applieddatamining.com