

ADAMS

Advanced **D**ata mining **A**nd **M**achine learning **S**ystem

Module: adams-weka



Peter Reutemann

December 22, 2016

©2009-2016



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato



Except where otherwise noted, this work is licensed under
<http://creativecommons.org/licenses/by-sa/4.0/>

Contents

1	Introduction	7
2	Flow	9
2.1	Conversions	9
2.2	Conditions	10
2.3	Actors	10
2.4	Templates	12
3	Classification and Regression	15
3.1	Basic	16
3.1.1	Loading data	16
3.1.2	Building models	17
3.1.3	Preprocessing	17
3.1.4	Evaluation	20
3.1.5	Making predictions	25
3.2	Advanced	26
3.2.1	Learning curves	26
3.2.2	Experiments	26
3.2.3	Optimization	26
3.2.4	Provenance	27
3.2.5	Partial Least Squares	28
4	Clustering	29
4.1	Building models	29
4.2	Evaluating clusterers	30
4.3	Clustering data	30
5	Attribute selection	33
6	Visualization	35
6.1	Preview browser	35
6.2	Instance Compare	37
6.3	Instance Explorer	38
6.4	Other visualizations	40
7	Tools	41
7.1	Append datasets	41
7.2	Arff viewer	42
7.3	Batch-filter datasets	43

7.4	BayesNet Editor	46
7.5	Dark Lord	47
7.6	Dataset compatibility	48
7.7	Explorer	49
7.8	Hermione	51
7.9	Investigator	52
7.10	Make compatible datasets	53
7.11	Merge datasets	55
7.12	Workbench	58
Bibliography		59

List of Figures

3.1	Flow for loading a local dataset.	16
3.2	The dataset that got loaded from disk.	16
3.3	Flow for generating and displaying an artificial dataset.	17
3.4	Flow for building J48 model on a dataset and outputting the model.	18
3.5	J48 model output.	18
3.6	Flow for comparing results generated from original and preprocessed “slug” data [4].	18
3.7	Evaluation summary on “slug” dataset (original).	19
3.8	Evaluation summary on “slug” dataset (log-transformed).	19
3.9	Classifier errors on “slug” dataset (original).	19
3.10	Classifier errors on “slug” dataset (log-transformed).	19
3.11	Cross-validating a classifier and outputting the summary.	20
3.12	Summary output of a cross-validated classifier.	20
3.13	Text file with command-lines of various classifiers.	21
3.14	Cross-validating classifier set ups read from a text file and displaying the evaluation summaries.	21
3.15	Summary outputs of cross-validated classifiers.	21
3.16	Flow for evaluating built classifier on a separate test set.	21
3.17	Summary output of classifier evaluated on separate test set.	21
3.18	Flow for building/evaluating classifier on a random split.	22
3.19	Summary output of classifier built/evaluated on random split.	22
3.20	Flow for evaluating classifier on separate train/test set.	23
3.21	Summary output of classifier evaluated on separate train/test set.	23
3.22	Flow for evaluating updateable classifier on data stream.	23
3.23	Summary output of classifier evaluated on data stream.	24
3.24	ROC curves of classifier evaluated on data stream.	24
3.25	Flow for displaying the “accumulated error” of a two classifiers.	24
3.26	The “accumulated error” of LinearRegression and GaussianProcesses.	24
3.27	Flow for classifying new data and outputting the class distributions.	25
3.28	The generated class distributions for the new data.	25
3.29	Flow for generating learning curve for incremental classifier.	26
3.30	Generated learning curve (incremental).	26
3.31	Flow for generating learning curve for batch classifier.	27
3.32	Generated learning curve (batch).	27
3.33	Flow for cross-validating a classifier on a pre-processed dataset.	27
3.34	Provenance display.	27

4.1	Building a clusterer and outputting the model.	29
4.2	Cluster model output.	29
4.3	Building a clusterer incrementally and outputting the model. . .	30
4.4	Cluster model outputs, generated every 25 instances.	30
4.5	Flow for clustering new data.	31
4.6	Generated cluster.	31
5.1	Flow for performing attribute selection (reduction).	33
5.2	Summary of the reduction.	34
5.3	The reduced dataset.	34
6.1	Default preview for classifier.	35
6.2	Preview of tree-generating classifier.	36
6.3	Preview for graph generated by BayesNet classifier.	36
6.4	Comparing two datasets.	37
6.5	Viewing data in the Instance explorer.	38
6.6	Viewing data in the Instance explorer.	39
6.7	Viewing data in the Instance explorer.	39
7.1	Appending datasets wizard.	41
7.2	WEKA's Arff viewer displaying the iris UCI dataset.	42
7.3	Selecting input files for batch-filtering.	43
7.4	Configuring the filter.	44
7.5	Selecting the output directory for the filtered datasets.	44
7.6	The dialog when batch-filtering was successful.	45
7.7	BayesNet editor displaying a net.	46
7.8	Dark Lord setup.	47
7.9	Dark Lord run.	47
7.10	Compatibility output for two datasets.	48
7.11	Explorer interface with menus.	49
7.12	Saving/restoring of workspaces.	50
7.13	Hermione setup.	51
7.14	Hermione run.	51
7.15	Weka Investigator.	52
7.16	Selecting input files for making compatible datasets.	53
7.17	Configuring the readers for the datasets.	54
7.18	Selecting the output directory for the compatible ARFF files. . .	54
7.19	The dialog when dataset generation was successful.	54
7.20	Selecting input files for merging.	55
7.21	Configuring the merge.	56
7.22	Selecting the output directory for the filtered datasets.	56
7.23	The dialog when merging was successful.	57
7.24	The merged dataset.	57
7.25	The Workbench interface.	58

Chapter 1

Introduction

The *adams-weka* module offers most of the functionality found in WEKA [2]: pre-processing, classification and regression, clustering, attribute selection, data visualization and visualization of results/models. But it does not stop there: the module also contains other features for optimization, experiment generation that are not available from WEKA, be it Explorer or KnowledgeFlow. It is assumed that you are familiar with WEKA¹ and machine learning in general, as common terms are not explained again.

If you have used WEKA's KnowledgeFlow before, then you will have to forget (mostly) everything that you know about setting up workflows. ADAMS does things quite differently in comparison to the WEKA. Additionally, ADAMS offers a range of general purpose actors that allow you to go further.

The manual is split into several sections, with: *classification and regression* and *clustering* comprising the most important sections.

¹If you haven't used WEKA before, check out the Data Mining book [3], which gives you a good introduction to machine learning, data mining and WEKA.

Chapter 2

Flow

The *adams-weka* module has a comprehensive set of actors and conversions that allow you to build powerful flows using WEKA's functionality. The following sections give a quick overview of available functionality. If you are interested in flow examples, check out chapters 3 and 4.

2.1 Conversions

This module offers additional schemes for the *Convert* transformer:

- *AdamsInstanceToWekaInstance* – converts an ADAMS instance into a WEKA one.
- *MatchWekaInstanceAgainstFileHeader* – uses a dataset header stored in a file to convert the string attributes of the instance passing through into nominal ones (and vice versa).
- *MatchWekaInstanceAgainstStorageHeader* – uses a dataset header obtained from storage to convert the string attributes of the instance passing through into nominal ones (and vice versa).
- *ReportToWekaInstance* – turns a *Report* object into a WEKA instance.
- *SpreadSheetToWekaInstances* – turns a spreadsheet object into a WEKA dataset.
- *WekaCapabilitiesToInstances* – generates a dataset with random data that adheres to the capabilities retrieved from a **CapabilitiesHandler**.
- *WekaCapabilitiesToSpreadSheet* – generates a spreadsheet with the capabilities retrieved from a **CapabilitiesHandler**.
- *WekaCommandToCode* – turns a WEKA command-line into code snippets.
- *WekaInstancesToSpreadSheet* – turns a WEKA dataset into a spreadsheet object.
- *WekaInstanceToAdamsInstance* – turns a WEKA instance into an ADAMS one.
- *WekaPredictionContainerToSpreadSheet* – generates a spreadsheet object from a prediction container (useful for display).

2.2 Conditions

The following boolean conditions, e.g., used in the *IfThenElse* or *Switch* control actors, are available:

- *AdamsInstanceCapabilities* – checks an ADAMS intance against the specified capabilities that it must satisfy.
- *WekaCapabilities* – checks a WEKA instance against the specified capabilities.
- *WekaClassification* (used in conjunction with *Switch*) – uses the returned classification index to determine which branch of the switch statement should be used; for all other control actors, the condition evaluates to “true” if an index is returned; condition works only with nominal classes.

2.3 Actors

The following sources are available:

- *WekaClassifierGenerator* – generates parameter sweeps for
- *WekaClassifierSetup* – outputs a single classifier setup.
- *WekaClustererGenerator* – generates parameter sweeps for
- *WekaClustererSetup* – outputs a single clusterer setup.
- *WekaDatabaseReader* – reads data from a database into WEKA’s internal format.
- *WekaDataGenerator* – generates artificial data using WEKA’s data generators.
- *WekaFilterGenerator* – generates parameter sweeps for filters.
- *WekaNewInstances* – simple source for generating empty datasets.
- *WekaSelectDataset* – source for selecting datasets interactively.

These transformers:

- *WekaAccumulatedError* – extracts all the errors collected during an evaluation, sorted according to magnitude and creates plot output, for comparing classifier performances (most useful for numeric classes).
- *WekaAggregatedEvaluations* – aggregates incoming Evaluation objects and forwards the current, aggregate state.
- *WekaAttributeIterator* – iterates through the names of a dataset and outputs them.
- *WekaTrainClassifier* – used for generating a trained model using a dataset.
- *WekaChooseAttributes* – allows the user to interactively select attributes to keep in a dataset.
- *WekaClassifierOptimizer* – applies a classifier optimizer (e.g., GridSearch or MultiSearch) to a dataset and then forwards the best (untrained) setup.
- *WekaClassifierRanker* – evaluates an array of classifier setups on a dataset and outputs the top X performing setups.
- *WekaClassifying* – uses a serialized (or callable) model to make predictions on incoming data.
- *WekaClassSelector* – sets the class attribute in a dataset.

- *WekaTrainClusterer* – trains a cluster algorithm setup on a dataset.
- *WekaClustering* – applies a serialized (or callable) model to incoming data.
- *WekaCrossValidationEvaluator* – performs cross-validation on an incoming dataset using a referenced classifier setup.
- *WekaCrossValidationSplit* – generates train/test set splits like cross-validation would generate.
- *WekaEvaluationSummary* – generates a summary for an Evaluation.
- *WekaEvaluationValuePicker* – retrieves a single statistic from an Evaluation.
- *WekaEvaluationValues* – generates a spreadsheet with the selected statistics from an Evaluation.
- *WekaExperiment* – executes a WEKA experiment, like in the Experimenter.
- *WekaExperimentEvaluation* – evaluates a WEKA experiment, generating text output of various sorts.
- *WekaExtractArray* – extracts a row or column from a WEKA dataset (using the internal format).
- *WekaFileReader* – reads any dataset that WEKA can handle, either outputs the header, the complete dataset or row-by-row.
- *WekaFilter* – applies a WEKA filter to the data.
- *WekaGetCapabilities* – retrieves the capabilities of a *CapabilitiesHandler* (e.g., Filter or Classifier).
- *WekaGetInstanceValue* – retrieves an attribute’s value from a dataset row.
- *WekaGetInstancesValue* – retrieves an attribute’s value from a dataset.
- *WekaInstanceBuffer* – buffers either incoming instance objects and outputs datasets or outputs instance objects when getting datasets.
- *WekaInstanceDumper* – for dumping dataset rows into files, one row at a time (ARFF or CSV).
- *WekaInstanceEvaluator* – adds an attribute with the value returned by an instance evaluator.
- *WekaInstanceFileReader* – outputs ADAMS instance objects.
- *WekaInstancesAppend* – creates one large dataset from multiple ones, by appending them one after the other.
- *WekaInstancesInfo* – outputs information on a dataset.
- *WekaInstancesMerge* – allows the merging of several datasets (side-by-side).
- *WekaInstanceStreamPlotGenerator* – generates plot containers from a range of attributes of instance objects passing through (i.e., you can plot several attributes in one go).
- *WekaModelReader* – reads a serialized model.
- *WekaMultiLabelSplitter* – splits a datasets with multiple class attributes (“multi-label”) into ones with only a single class attribute.
- *WekaNewInstance* – creates an instance object with only missing values using a dataset as template.
- *WekaPredictionsToInstances* – turns WEKA predictions into a WEKA dataset (actual, predicted, etc).
- *WekaPredictionsToSpreadSheet* – turns WEKA predictions into a spread-

sheet (actual, predicted, etc).

- *WekaRandomSplit* – generates a random split of a dataset.
- *WekaRegexToRange* – generates a range string using a regular expression applied to the names of a dataset.
- *WekaRelationName* – simply outputs the name of the dataset.
- *WekaRenameRelation* – renames a dataset.
- *WekaReorderAttributesToReference* – reorders the attributes in incoming Instance/Instances based on an order defined in a reference dataset (callable actor or file).
- *WekaSetInstanceValue* – sets a specific attribute value in an instance object.
- *WekaSetInstancesValue* – sets a specific attribute value in a dataset object.
- *WekaStoreInstance* – appends the passing through instance
- *WekaStreamFilter* – works the same as *WekaFilter* but only allows stream filters to be selected.
- *WekaSubsets* – splits dataset into subsets using the unique values of an attribute to identify subsets.
- *WekaTestSetEvaluator* – evaluates a trained classifier on a dataset obtained from a callable actor.
- *WekaTextDictionaryReader* – reads in a directory with the documents in the sub-directories representing different classes.
- *WekaTrainTestSetEvaluator* – evaluates a referenced classifier using the incoming train/test split.

And these sinks:

- *WekaClassifierErrors* – displays the errors of a classifier.
- *WekaCostCurve* – generates a cost curve.
- *WekaDatabaseWriter* – writes a dataset to a database.
- *WekaExperimentGenerator* – generates a WEKA experiment by adding the incoming classifier setups and writing it to disk.
- *WekaFileWriter* – writes a dataset to any file format that WEKA can handle.
- *WekaInstancesDisplay* – displays datasets in table format.
- *WekaInstancesPlot* – plots one attribute vs another.
- *WekaInstanceViewer* – visualizes incoming WEKA or ADAMS instance objects the same way as the *Instance Explorer* tool does.
- *WekaModelWriter* – writes a model container or classifier/clusterer to disk.
- *WekaThresholdCurve* – displays threshold curves like, receiver-operator curve (ROC) or precision/recall.

2.4 Templates

Here are some templates that make the flow development for WEKA easier:

- *InstanceDumperVariable* – generates a variable for the *WekaInstanceDumper* actor which contains an ARFF/CSV filename prefix aligned with

the flow's filename, i.e., the ARFF/CSV file will always get placed in the same location as the flow.

Chapter 3

Classification and Regression

WEKA's main strength lies in its large number of classification and regression schemes. Most of the documentation will cover this functionality therefore.

We start out with some basic WEKA functionality, like loading and preprocessing data, building models and evaluating them. That includes visualization of the results and models as well. After that we will cover more advanced features like learning curves, experiment generation and evaluation, optimization of classifiers and also the current provenance support in ADAMS.

3.1 Basic

In this section we describe how to perform basic WEKA functionality that you are used to perform with the Explorer, but in the workflow context. Instead of having to repeat the same steps, like loading and preprocessing data, whenever you update your data, a flow allows you to define the steps apriori and then merely re-execute them time and time again. Also, flows make it very easy to *document* all the steps that you perform, not just merely recording what you are doing.

3.1.1 Loading data

Before we can build any models, we have to have data at hand, of course. So the first step will be to obtain data from somewhere, whether that is by loading a local dataset or by downloading a remote dataset.

To start, we will be loading files that are stored locally. The actor used for loading datasets is the *WekaFileReader* transformer. This actor does not have an option for the file to load. Instead, it expects a file name, string or URL object to arrive at its input port. In order to supply a local file, we use the *SingleFileSupplier* source, which allows us to specify a single file that gets forwarded in the flow. If required, one can also use the *MultiFileSupplier* or *DirectoryLister* sources¹, which can forward multiple file names instead of just one. The latter one is especially handy, if the files are not known in advance, e.g., generated on the fly. In order to display the loaded data, we use the *WekaInstancesDisplay* sink actor, which displays the data in a nice tabular format. Figure 3.1 shows the flow for loading the dataset and Figure 3.2 the generated output.




Figure 3.1: Flow for loading a local dataset.

1: family	2: product-type	3: steel	4: carbon	5: hardness	6: temper	7: rolling	8: condition	9: fatigue life
1	C	A	0.0	0.0	0.0	0.0	0.0	0.0
2	C	R	0.0	0.0	0.0	0.0	0.0	0.0
3	C	R	0.0	0.0	0.0	0.0	0.0	0.0
4	C	A	0.0	60.0	0.0	0.0	0.0	0.0
5	C	A	0.0	60.0	0.0	0.0	0.0	0.0
6	C	A	0.0	45.0	0.0	0.0	0.0	0.0
7	C	R	0.0	0.0	0.0	0.0	0.0	0.0
8	C	R	0.0	0.0	0.0	0.0	0.0	0.0
9	C	R	0.0	0.0	0.0	0.0	0.0	0.0
10	C	A	0.0	0.0	0.0	0.0	0.0	0.0
11	C	R	0.0	0.0	0.0	0.0	0.0	0.0
12	C	R	0.0	0.0	0.0	0.0	0.0	0.0
13	C	R	0.0	0.0	0.0	0.0	0.0	0.0
14	C	A	0.0	45.0	0.0	0.0	0.0	0.0
15	C	A	0.0	0.0	0.0	0.0	0.0	0.0
16	C	A	0.0	0.0	0.0	0.0	0.0	0.0
17	C	A	10.0	0.0	0.0	0.0	0.0	0.0
18	C	A	0.0	80.0	0.0	0.0	0.0	0.0
19	C	A	0.0	0.0	0.0	0.0	0.0	0.0
20	C	A	0.0	70.0	0.0	0.0	0.0	0.0
21	C	A	0.0	0.0	0.0	0.0	0.0	0.0
22	C	K	55.0	0.0	0.0	0.0	0.0	0.0
23	C	A	0.0	85.0	0.0	0.0	0.0	0.0

Figure 3.2: The dataset that got loaded from disk.

In this example² we let the *WekaFileReader* determine the correct file loader automatically, based on the file extension. If this automatic determination should fail, you can always check the “useCustomLoader” checkbox and then configure the appropriate loader yourself.

Another feature of this actor is the ability to output the dataset row by row (option “incremental”). This is very handy in case of very large files, where loading into memory could pose a problem. Even though the incremental feature

¹adams-weka-crossvalidate_classifier_multiple_datasets.flow

²adams-weka-load_dataset.flow

works for any file type that WEKA can read, truly incremental, i.e., memory-efficient, loading is only possible if the underlying loader also supports incremental loading. In any other case, the dataset gets loaded fully into memory before being forwarded row by row.

Finally, artificial data can be generated within ADAMS as well. Using the *WekaDataGenerator* source, any WEKA data generator can be used to output data. The flow³ depicted in Figure 3.3 generates a small dataset using the “Agrawal” data generator.




Figure 3.3: Flow for generating and displaying an artificial dataset.

3.1.2 Building models

After having sorted out the loading of the data, it is time to check out how to build models. Since we are using supervised algorithms, we have to make sure that the datasets have a class attribute set. The *WekaClassSelector* actor allows the setting of the class attribute, in the default setting it simply uses the last attribute as the class attribute. With the *WekaTrainClassifier* actor you can choose a callable classifier to be built. You define a callable classifier by adding a *WekaClassifierSetup* source to the *CallableActors* standalone, which you then reference in your *WekaTrainClassifier* actor. By default, the *WekaTrainClassifier* actor outputs a container that comprises the built model and the header of the training set. In order to extract either of the container items, you need to use the *ContainerValuePicker* control actor. Figure 3.4 demonstrates how to train a J48 classifier on dataset and then displaying the built model (see Figure 3.5)⁴.

A built model can be saved to disk (and then re-used later) using the *WekaModelWriter*. The file generated can also be loaded in the WEKA Explorer again and applied to another test set there⁵.

3.1.3 Preprocessing

A very important, but often underrated step is preprocessing. Unless your data is properly cleaned up and in the right format, your models will not be very meaningful. Preprocessing steps can be done within the flow using the *WekaFilter* transformer, which wraps around a single WEKA filter. One either

³adams-weka-data_generator.flow

⁴adams-weka-build_classifier-output_only_model.flow

⁵adams-weka-build_classifier-save_model.flow

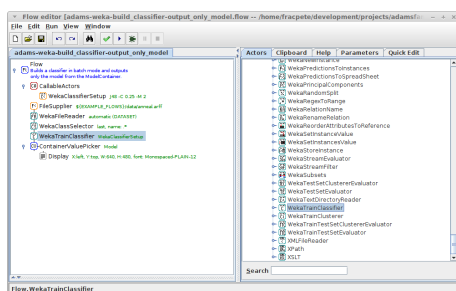


Figure 3.4: Flow for building J48 model on a dataset and outputting the model.

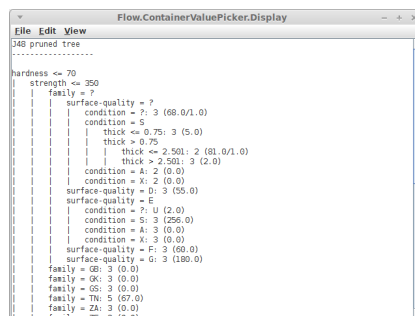


Figure 3.5: J48 model output.

chains multiple actors together or uses the *weka.filters.MultiFilter* meta-filter to executed several filter sequentially in a single actor.

In Figure 3.6 we are investigating the impact of preprocessing on the “slug” dataset [4]. The flow⁶ cross-validates *LinearRegression* on the original and log-transformed data. The log-transformed data is generated by applying the *AdExpression* filter on each of the two attributes of the dataset and then deleting the original ones. In each case, original or preprocessed, it displays the evaluation summary and classifier errors.




Figure 3.6: Flow for comparing results generated from original and preprocessed “slug” data [4].

Figures 3.7 and 3.8 show the evaluation summary, for the original and the log-transformed data. The log-transformed dataset gets not only a better correlation coefficient, but also smaller errors.

Figures 3.9 and 3.10 display the classifier errors. It is obvious from the funny log-shaped curve, that *LinearRegression* built on the original data is not a very good model. Something that is not so obvious by just looking at the correlation coefficient: 0.9056 is not bad.

⁶adams-weka-filter.data.flow




Figure 3.7: Evaluation summary on “slug” dataset (original).

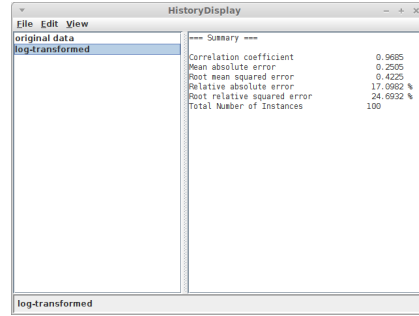


Figure 3.8: Evaluation summary on “slug” dataset (log-transformed).

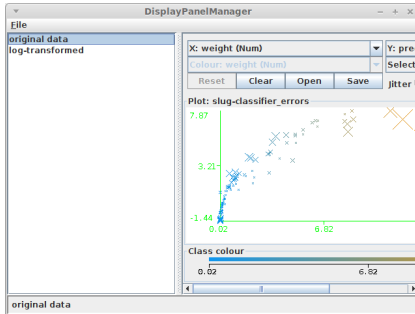


Figure 3.9: Classifier errors on “slug” dataset (original).

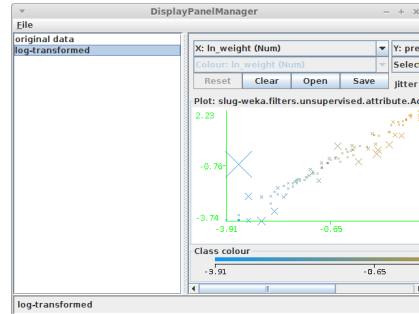


Figure 3.10: Classifier errors on “slug” dataset (log-transformed).

This flow can be quickly extended to accommodate other preprocessing techniques, all very easily comparable in the graphical output.

In this example the preprocessing was rather specific. On the other hand, if you are working mainly in a particular data domain, like spectral analysis of some kind, then certain preprocessing steps will always be same. In this case, it makes sense to store these externally in a *preprocessing library* which you then link to using external actors (see manual for the *adams-core* module for more details). This reduces duplication and you will only have to update the preprocessing step in a single location.

Instead of batch-filtering data, you can also filter streams of *weka.core.Instance* objects, using the *WekaStreamFilter* transformer. This filter offers a subset of WEKA’s filters, which don’t need a batch of data to be initialized with before being able to process data.


```

weka.classifiers.trees.J48 -C 0.25 -M 2
weka.classifiers.meta.FilteredClassifier -F
  "weka.filters.unsupervised.attribute.Remove -R 1-3" -W weka.classifiers.trees.J48
  -- -C 0.25 -M 2
weka.classifiers.functions.SMO -C 1.0 -L 0.0010 -P 1.0E-12 -N 0 -M -V -1 -W 1 -K
  "weka.classifiers.functions.supportVector.RBFSKernel -C 250007 -G 0.01"

```

Figure 3.13: Text file with command-lines of various classifiers.




Figure 3.14: Cross-validating classifier set ups read from a text file and displaying the evaluation summaries.

Flow.evaluate.classifier.HistoryDisplay

weka.classifiers.trees.J48 -C 0.25 -M 2
weka.classifiers.meta.FilteredClassifier
weka.classifiers.functions.SMO -C 1.0 -L 0.0010 -P 1.0E-12 -N 0 -M -V -1 -W 1 -K
"weka.classifiers.functions.supportVector.RBFSKernel -C 250007 -G 0.01"

Summary

Correctly Classified Instances 854 95.1002 %
Incorrectly Classified Instances 48 4.8998 %
Kappa statistic 0.9264
Mean absolute error 0.1125
Root mean squared error 19.8856 %
Relative absolute error 43.5788 %
Root relative squared error 66.8573 %
Coverage of cases (0.95 level) 20.2116 %
Total Number of Instances 890

Detailed Accuracy By Class

TP Rate	FP Rate	Precision	Recall	F-Measure	ROC Area
0.875	0.004	0.938	0.875	0.737	0.744
0.848	0.011	0.903	0.848	0.875	0.863
0.878	0.121	0.963	0.878	0.970	0.954
0.900	0.000	0.900	0.900	0.900	1.000
1.000	0.000	1.000	1.000	1.000	1.000
0.875	0.000	0.844	0.875	0.790	0.745
0.952	0.004	0.952	0.952	0.950	0.974

Weighted Avg

Figure 3.15: Summary outputs of cross-validated classifiers.

Test set evaluation

Simply testing a built classifier on a test set is useful when you are always intending to save the generated model to a file, but also want to keep an eye on the performance. In this case, you can very easily extend your current flow for building and saving the model. First, add a callable actor that loads the separate training set from disk. Second, add a *Tee* control actor that performs the evaluation using the *WekaTestSetEvaluator* and *WekaEvaluationSummary* transformers and a *Display* sink for showing the results⁹. The full flow and the generated output are shown in Figures 3.16 and 3.17.




Figure 3.16: Flow for evaluating built classifier on a separate test set.

Flow.evaluate.Display

Summary

Correctly Classified Instances 302 99.0164 %
Incorrectly Classified Instances 3 0.9836 %
Kappa statistic 0.9747
Mean absolute error 0.0047
Root mean squared error 0.0562
Relative absolute error 3.6361 %
Root relative squared error 27.235 %
Coverage of cases (0.95 level) 99.0164 %
Mean rel. region size (0.95 level) 19.6667 %
Total Number of Instances 305

Figure 3.17: Summary output of classifier evaluated on separate test set.

⁹adams-weka-build_classifier_evaluate_on_testset.flow

Train/test set evaluation

An evaluation using separate train and test set can be used, if you don't want to keep the evaluated model, but you are only interested in the evaluation output. The evaluation actor in this case is the *WekaTrainTestSetEvaluator* transformer. This actor accepts *WekaTrainTestSetContainer* data tokens. To generate this container you have several options:

- *WekaRandomSplit* – splits a single dataset into a train and test set, based on the percentage supplied by the user.
- *WekaCrossValidationSplit* – Generates train/test splits like they occur in cross-validation. Useful, if you want to inspect the various models built during cross-validation, not just the summary.
- *MakeContainer* – manually generating a container from two individually loaded datasets.

Figures 3.18 and 3.19 show how to use the *RandomSplit* actor in the evaluation process¹⁰. For simulating cross-validation, simply exchange the *WekaRandomSplit* actor with a *WekaCrossValidationSplit* one (you might also want to change from *Display* to *HistoryDisplay*, to keep better track of the various evaluations).

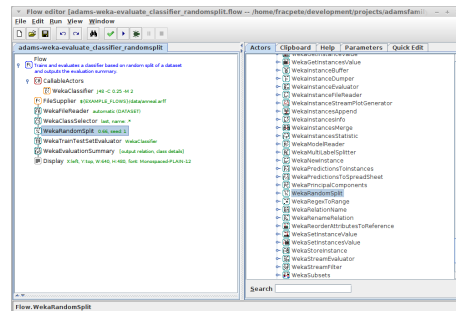


Figure 3.18: Flow for building/evaluating classifier on a random split.

Flow.Display

Relation: adams

==== Summary ====

Correctly Classified Instances	294	96.3334 %
Incorrectly Classified Instances	11	3.6666 %
Kappa statistic	0.9316	
Mean absolute error	0.0227	
Root mean squared error	0.1503	
Relative absolute error	9.439 %	
Root relative squared error	39.7443 %	
Coverage of cases (0.95 level)	96.6885 %	
Mean rel. region size (0.95 level)	17.9762 %	
Total Number of Instances	305	

==== Detailed Accuracy By Class ====

	TP Rate	FP Rate	Precision	Recall	F-Measure	MCC	ROC Area	PRC Area	Class
0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.987	0.270	1
0.999	0.029	0.795	0.999	0.973	0.982	0.980	0.910	0.910	2
0.978	0.027	0.991	0.978	0.985	0.982	0.982	0.989	0.989	3
0.000	0.000	0.000	0.000	0.000	0.000	0.000	?	?	4
1.000	0.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	5
1.000	0.003	0.929	1.000	0.993	0.992	0.998	0.929	0.929	U
Weighted Avg.	0.964	0.104	0.952	0.964	0.957	0.920	0.982	0.967	

Figure 3.19: Summary output of classifier built/evaluated on random split.

Figures 3.20 and 3.21 display the flow¹¹ for manually creating a container using the general purpose *MakeContainer* source actor. In order to assemble a container, you need to know **what** type of container you want to create (the type is normally listed in the “Help” of an actor), **where** to obtain the data from (i.e., the callable actors) and **how** to store the data (i.e., under which name in the container).

Stream evaluation

Though WEKA is usually used for batch-training, it is also possible to perform incremental training and evaluation, as long as the classifier in use is an *updateable* one (i.e., it implements the *weka.classifiers.Updateable* interface). In such a scenario, you only have to read in the data incrementally (changing the set up in the *WekaFileReader* to *INCREMENTAL*) and use the *WekaStreamEvaluator*

¹⁰adams-weka-evaluate_classifier_randomsplit.flow

¹¹adams-weka-assemble_trainittestset_container_and_evaluate_classifier.flow




Figure 3.20: Flow for evaluating classifier on separate train/test set.

Summary		
Correctly Classified Instances	362	99.0164 %
Incorrectly Classified Instances	3	0.9836 %
Kappa statistic	0.9747	
Mean absolute error	0.0047	
Root mean squared error	0.0562	
Relative absolute error	3.5448 %	
Root relative squared error	22.2596 %	
Coverage of cases (0.95 level)	99.0164 %	
Mean rel. region size (0.95 level)	16.6867 %	
Total Number of Instances	365	

Figure 3.21: Summary output of classifier evaluated on separate train/test set.

transformer for performing the evaluation. The *WekaStreamEvaluator* actor use prequential evaluation, i.e., first evaluate, then train. Figure 3.22 shows a flow that evaluates a *NaiveBayesUpdateable* on a stream of data, displaying the textual evaluation summary and ROC (receiver operator curve) every 100 instances that come through.

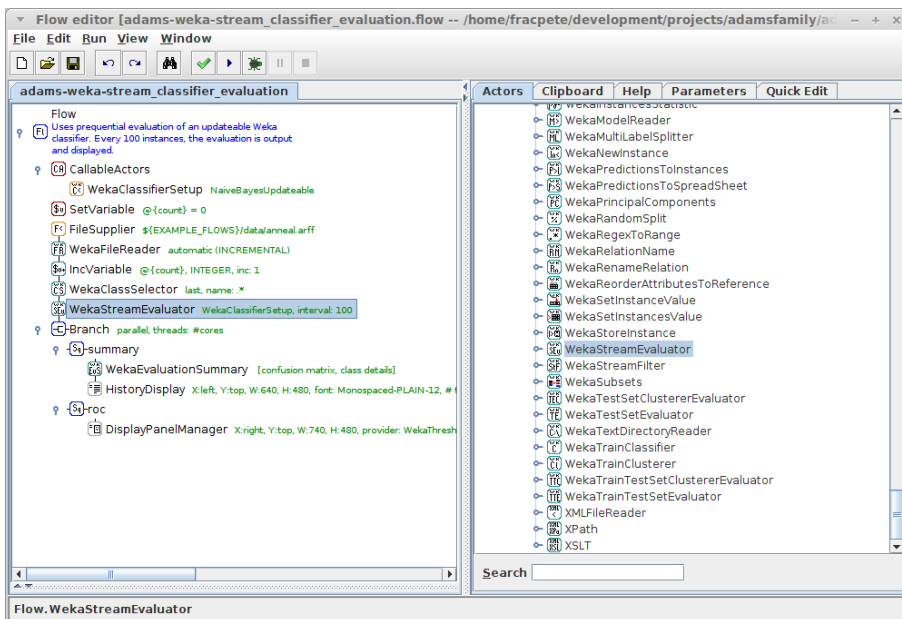


Figure 3.22: Flow for evaluating up-dateable classifier on data stream.

Visualization

You have already encountered the display of the classifier errors (in Figure 3.9). The sink for displaying these errors is *WekaClassifierErrors*, which takes an *Evaluation* object as input. If you want to evaluate and display multiple classifiers then you have to use the *DisplayPanelManager* with the *WekaClas-*




Figure 3.23: Summary output of classifier evaluated on data stream.




Figure 3.24: ROC curves of classifier evaluated on data stream.

sifierErrors actor as “panelProvider”. The *DisplayPanelManager* actor offers a history of generated panels, like the *HistoryDisplay* does for plain text.

Another interesting visualization is the *WekaAccumulatedError* transformer. This transformer takes also an *Evaluation* object and then turns it into a special sequence of plot containers: it creates a sequence of the prediction errors that were obtained during an evaluation and outputs them sorted, from smallest to largest¹². The Figures 3.25 and 3.26 show the flow and the generated output respectively. As you can see from the graph, GaussianProcesses generates consistently larger errors than LinearRegression, which only seems to have a few big outliers (steep increase at the end).

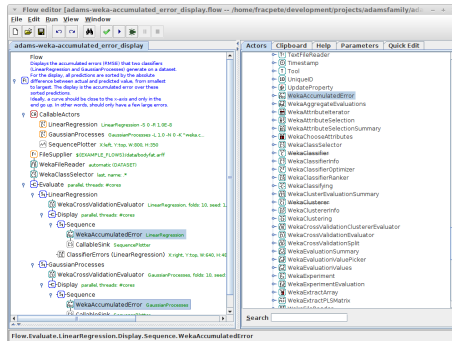


Figure 3.25: Flow for displaying the “accumulated error” of a two classifiers.




Figure 3.26: The “accumulated error” of LinearRegression and GaussianProcesses.

¹²adams-weka-accumulated_error_display.flow

3.1.5 Making predictions

Of course, building models is only part of the picture. You will want to use this model as well and make predictions with it. The actor for making predictions on incoming data (i.e., single instance objects) is the *WekaClassifying* actor. This actor can either use a serialized model or a callable actor that generates a trained classifier. The flow¹³ in Figure 3.27 uses the callable actor approach, training a classifier on a training set and then performing classifications on a test set, with the class distributions shown on screen (see Figure 3.28).

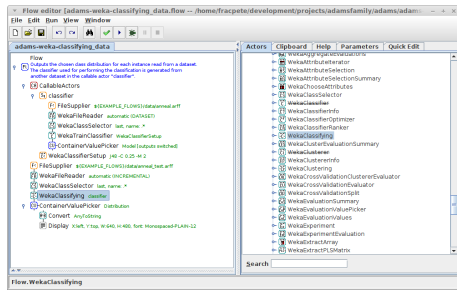


Figure 3.27: Flow for classifying new data and outputting the class distributions.




Figure 3.28: The generated class distributions for the new data.

¹³adams-weka-classifying_data.flow

3.2 Advanced

3.2.1 Learning curves

Classifiers are susceptible to the order and amount of data that they are trained with. Using learning curves, one can investigate how the data influences the classifier performance.

Figures 3.29 and 3.30 show a flow and the generated output of an incremental NaiveBayes classifier, with the classifier being evaluated against a test set every 10 training instances (*ConditionalTee* in conjunction with the *Counting* condition).¹⁴

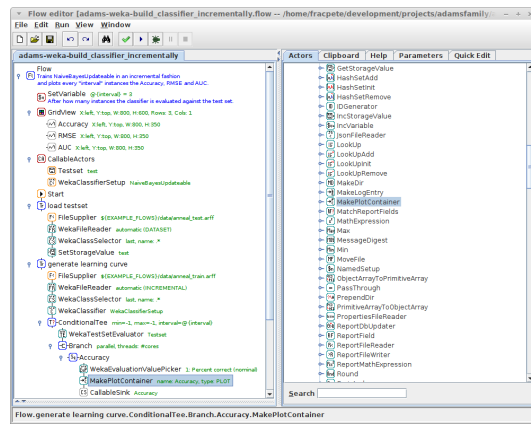


Figure 3.29: Flow for generating learning curve for incremental classifier.




Figure 3.30: Generated learning curve (incremental).

Incremental classifiers are great for generating these kind of graphs. But even for batch classifiers you can generate learning curves. Using the *WekaInstanceBuffer* transformer, it is possible to buffer instances as they come through and output datasets with which the batch classifier can get trained (and evaluated). Figures 3.31 and 3.30 show how to generate a learning curve for the decision tree classifier J48, being evaluated every 10 instances.¹⁵

3.2.2 Experiments

experiment generation¹⁶, execution and evaluation¹⁷

3.2.3 Optimization

setup generators¹⁸, ranker¹⁹, optimizer²⁰

¹⁴adams-weka-build_classifier_incrementally.flow

¹⁵adams-weka-classifier_learning_curve.flow

¹⁶adams-weka-experiment_generation.flow

¹⁷adams-weka-experiment.flow

¹⁸adams-weka-classifier_setup_generation.flow

¹⁹adams-weka-classifier_setup_ranking.flow

²⁰adams-weka-classifier_optimizer.flow




Figure 3.32: Generated learning curve (batch).

Machine learning related actors can keep track of what operations happened along the way, or in other words, provenance.

Enabled=true

The screenshot shows the Flow editor interface. The top bar includes a file explorer and a menu bar with 'File', 'Edit', 'Run', 'View', and 'Window'. The main canvas displays a flow diagram with several widgets. The 'Flow.PromoteDisplay' widget is selected, and its properties are shown in the right sidebar. The left sidebar shows the 'Flow' palette with various widgets like 'Flow.PromoteDisplay', 'Flow.PromoteDisplay', 'Flow.PromoteDisplay', etc. The right sidebar shows the 'Properties' panel for the selected widget, displaying various configuration options like 'Display', 'Display', 'Display', etc.




Figure 3.34: Provenance display.

²¹adams-weka-crossvalidate_classifier-display_provenance.flow

3.2.5 Partial Least Squares

Using the *WekaExtractPLSMatrix* transformer, you can extract various PLS matrices from a *PLSFilterWithLoadings* filter or a *PLSClassifierWeightedWithLoadings* classifier (or a *WekaModelContainer*, if this container should have a *PLSClassifierWeightedWithLoadings* classifier stored).²²

²²`adams-weka-extract-pls-matrix.flow`

Chapter 4

Clustering

Clustering behaves very much like Classification/Regression, the only difference being that it is an unsupervised learning process. This means that the flows won't contain a *WekaClassSelector* actor to set the class attribute in the loaded data. Due to the similarity, the section here will cover only the basics of clustering.

4.1 Building models

Building clustering models is as easy as building classification/regression models. Instead of the *WekaTrainClassifier* transformer, you use the *WekaTrainClusterer* one. Similar, you use a *WekaClustererSetup* source instead of the *WekaClassifierSetup* one to define (and output) a clusterer setup, placed inside a *CallableActors* standalone.

Figures 4.1 and 4.2 show a flow ¹ that builds a *SimpleKMeans* clusterer on a dataset (the class attribute gets removed using a *WekaFilter* actor) and the generated model gets displayed.

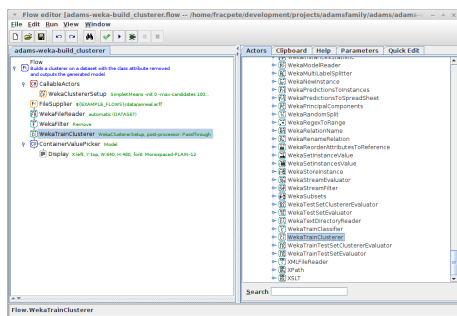


Figure 4.1: Building a clusterer and outputting the model.

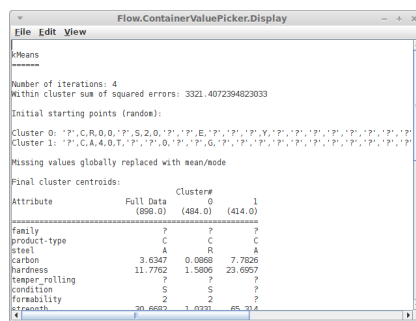


Figure 4.2: Cluster model output.

If the base cluster algorithm is an incremental one, i.e., one that implements the *weka.clusterers.UpdateableClusterer* interface, you can build your clustering

¹adams-weka-build_clusterer.flow

model incrementally as well. The flow ² in Figure 4.3 builds the CobWeb cluster algorithm incrementally and outputs the generated models every 25 instances (see Figure 4.4).




Figure 4.3: Building a clusterer incrementally and outputting the model.




Figure 4.4: Cluster model outputs, generated every 25 instances.

4.2 Evaluating clusterers

ADAMS offers transformers for evaluation clusterers on data similar to the ones for classification:

- *WekaClusterAssignments* - outputs the cluster assignments from the last evaluation (if available).
- *WekaClusterEvaluationSummary* - generates a string representation of a cluster evaluation (or container)
- *WekaCrossValidationClustererEvaluator* - cross-validates a clusterer on a dataset, generates log-likelihood.
- *WekaTestSetClustererEvaluator* - evaluates a built clusterer on a test set.
- *WekaTrainTestSetClustererEvaluator* - builds and evaluates a clusterer on the training and test set from a train/test-set container.

4.3 Clustering data

Clustering new data is done using the *WekaClustering* transformer, which takes a single instance as input and outputs the generated clustering information in form of a container (*WekaClusteringContainer*). You can either specify a serialized clusterer model to use or a callable actor to obtain the clusterer from. The flow ³ in Figure 4.5 shows how to build a clusterer and use it to cluster new data, outputting the cluster distributions (see Figure 4.6 for the generated output).

²adams-weka-build_clusterer_incrementally.flow

³adams-weka-clustering_data.flow

4.3. CLUSTERING DATA

31

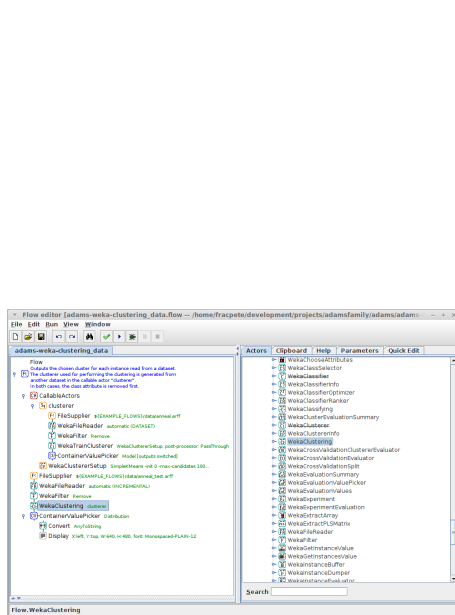


Figure 4.5: Flow for clustering new data.



Figure 4.6: Generated cluster.

Attribute selection

- *WekaAttributeSelection* – performs the attribute selection/ranking.
- *WekaAttributeSelectionSummary* – generates a summary from a attribute selection step.

The *WekaAttributeSelection* transformer outputs a container which can contain the following elements:

- ¹
- adams-weka-attribute-selection-subset.flow

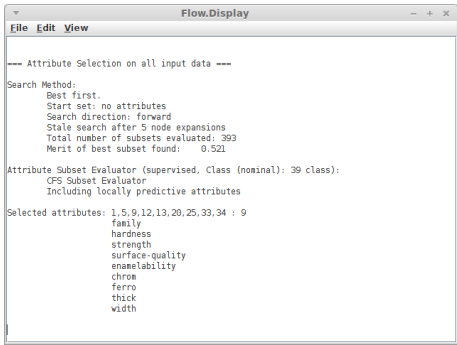


Figure 5.2: Summary of the reduction.

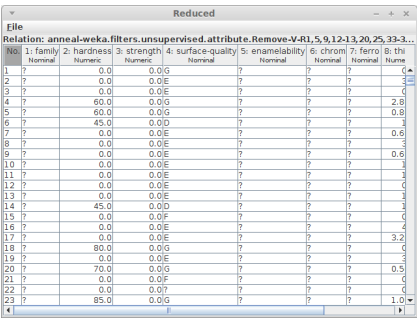


Figure 5.3: The reduced dataset.

- *Seed* – the seed value in case of cross-validation.
- *Folds* – the number of folds used in case of cross-validation.

Chapter 6

Visualization

6.1 Preview browser

The WEKA module comes with custom viewers for serialized files. Apart from the default view (Figure 6.1), you can also view the trees (Figure 6.2) and graphs (Figure 6.3) that some classifiers generate.




Figure 6.1: Default preview for classifier.




Figure 6.2: Preview of tree-generating classifier.

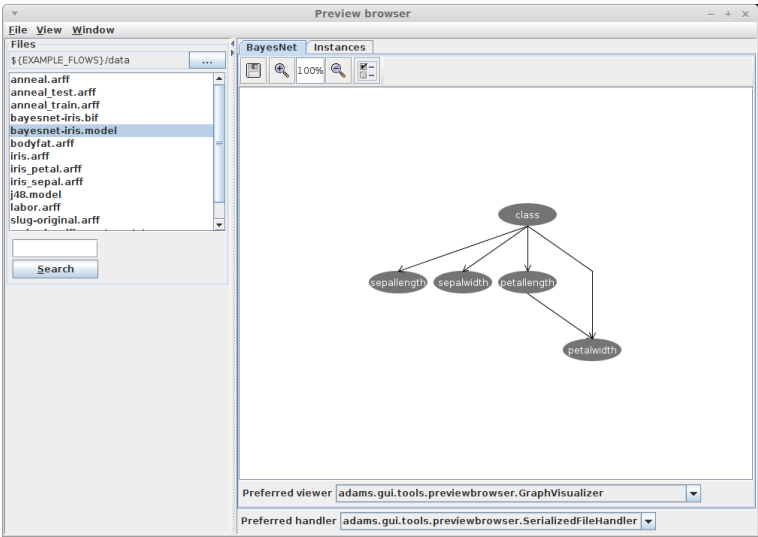


Figure 6.3: Preview for graph generated by BayesNet classifier.

6.2 Instance Compare

Quite often, you generate data with different or tweaked pre-processing techniques and you wonder how different the generated data looks like. The *Instance Compare* visualization allows you to graphically compare two datasets. You can either compare them row by row, or using a common attribute that can be used as unique row identifier.

Figure 6.4 shows a comparison of two datasets. Not only are the two rows overlaid, you also see the absolute difference plotter and a the correlation coefficient of the two being calculated.

If you don't want to compare all the attributes, you can restrict it to a subset, by using the *Att. range* text field. “first”, “second”, “third”, “last”, “last_1” (last minus 1) and “last_2” (last minus 2) are accepted indices. All other indices must be 1-based.

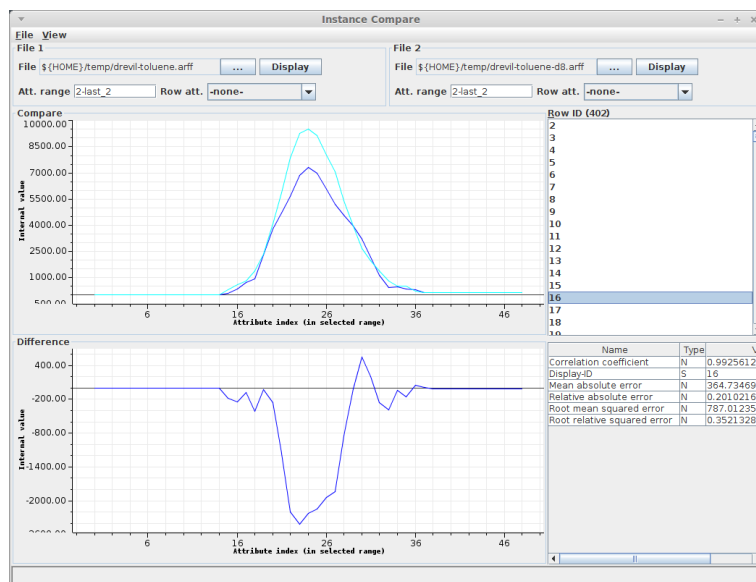


Figure 6.4: Comparing two datasets.

6.3 Instance Explorer

When generating datasets, it pays to check the generated output in multiple ways. For instance, whether the data rows generated are actually aligning properly. The *Instance Explorer* allows you to select a range of rows and columns from a dataset (see Figures 6.5 and 6.6), which are then displayed in a single graph.

Figure 6.7 shows a subset of the UCI dataset *waveform-5000*. The top graph of the two is a zoom into the full graph, with the bottom graph showing the area that was zoomed into.

Index	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	x11	x12	x13	x14
1	-0.23	-1.21	1.2	1.23	-0.1	0.12	2.49	1.19	1.34	0.58	1.22	2.3	4.65	5.82
2	0.38	0.38	-0.31	-0.09	1.52	1.35	1.49	3.81	2.33	1.34	1.45	3.7	3.08	5.01
3	-0.69	1	1.08	1.48	2.44	3.39	3.09	4.08	5.48	3.61	0.47	1.68	2.35	-1.04
4	0.4	0.68	0.27	1.39	1.03	-0.32	-1.23	-0.5	0.11	0.87	1.27	4.41	3.51	4.88
5	-0.81	1.59	-0.69	1.16	4.22	4.98	4.52	2.54	5.6	4.66	4.25	1.58	2.51	2.4
6	0.59	0.77	-0.61	1	1.8	2.08	2.16	3.59	4.08	3.63	4.27	4.43	3.45	2.47
7	-0.15	0.13	2.27	2.39	4	6.14	5.36	4.08	3.81	3.89	2.46	1.78	-1.43	0.7
8	-0.3	-0.42	0.25	-0.61	-1.39	-0.6	1.71	4.01	2.96	5.81	6.56	5.69	5.46	2.72
9	-1.45	2.71	3.04	3.21	4.26	5.01	6.24	5.09	3.95	4.84	2.15	-0.3	1.53	-1.19
10	0.28	0.97	-1.01	-2.34	-1.89	0.54	0.05	2.05	2.38	3.66	3.09	5.12	4.14	3.7
11	-1.09	-0.44	1.15	0.17	2.1	3.77	2.4	5.16	5.13	3.66	2.42	2.83	1.02	1.26
12	0.5	-1.23	-0.09	0.31	2.22	-0.02	2.24	-0.07	1.98	2.27	4.94	3.97	3.85	3.19
13	-0.23	-0.44	1.04	0.38	0.53	-0.95	0.97	1.57	1.48	4.88	4.7	4.73	5.38	4.65
14	-1.2	-1.04	-0.06	1.24	-1.41	0.3	1.23	3.2	3.53	3.25	5.93	4.94	5.79	3.47
15	-0.53	0.8	0.1	0.25	0.59	-0.38	-0.13	0.31	0.66	2.79	3.31	3.46	4.06	4.04
16	-0.17	-0.27	-0.46	2.09	-1.32	0.88	2.2	2.13	2.04	3.54	4.04	5.12	5.06	2.54
17	0.58	1.69	-0.77	0.36	0.37	0.49	0.87	1.67	1.64	2.29	5.7	5.64	4.2	3.62
18	-1.06	-1.26	-0.87	0.83	-0.38	-0.35	1	1.38	1.59	3.01	4.67	3.71	2.54	4.71
19	1.86	1.31	1.27	1.55	2.59	2.82	3.78	4.89	4.27	3.64	5.6	3.2	1.04	0.03
20	0.15	1.01	-0.18	0.68	1.61	1.83	1.66	2.22	2.9	2.8	1.62	2.96	3.07	3.33
21	0.25	-1.42	-1.39	2.87	0.56	2.03	1.12	1.28	1.97	5.59	6.16	5.66	3.07	3.31
22	0.49	0.08	0.14	0.45	2.66	2.34	2.55	0.44	1.67	1.63	1.66	1.53	2.82	3.33
23	-0.53	2.45	-0.22	3.77	2.97	2.51	4.76	4.51	3.03	2.23	1.52	2.16	1.35	1.66

Figure 6.5: Viewing data in the Instance explorer.

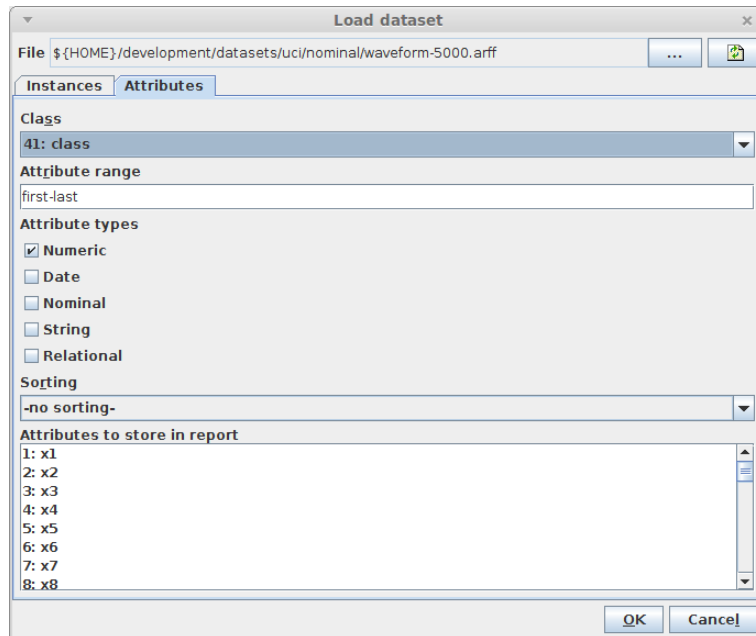


Figure 6.6: Viewing data in the Instance explorer.

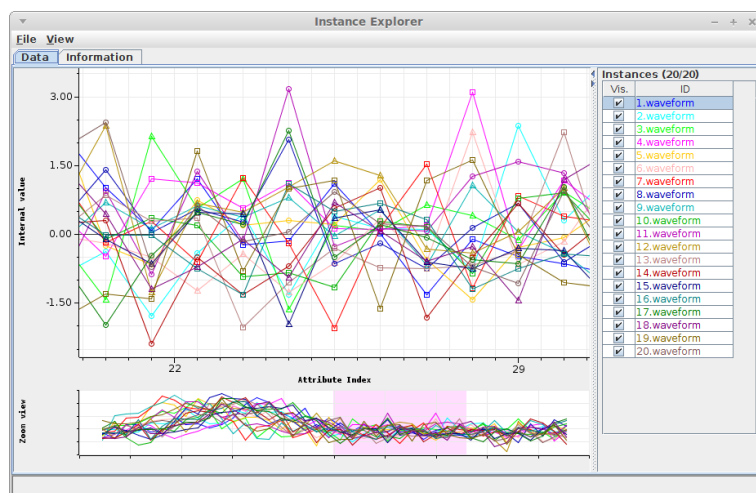


Figure 6.7: Viewing data in the Instance explorer.

6.4 Other visualizations

Below are some more visualizations, which are available through WEKA and have been added to the main menu:

- *Boundary visualizer* – visualizes classification boundaries
- *Cost curve* – for re-visualizing a previously displayed cost curve
- *Graph visualizer* – for displaying BayesNet graphs
- *Instances plot* – plots a attributes from a dataset against each other
- *Margin curve* – for re-visualizing a previously displayed margin curve
- *Plot attribute vs attribute* – allows to plot attributes against each other, using two user-selected subsets.
- *ROC* – for re-visualizing a previously displayed receiver operator curve
- *Tree visualizer* – for displaying trees graphs (digraph format¹)

¹https://en.wikipedia.org/wiki/DOT_%28graph_description_language%29

Chapter 7

Tools

7.1 Append datasets

If you have datasets with the same structure, i.e., exactly the same attributes, then you can use the *Append datasets* tool to combine them into a single, large dataset.

The screenshot in Figure 7.10 shows the wizard interface that guides you through the process.

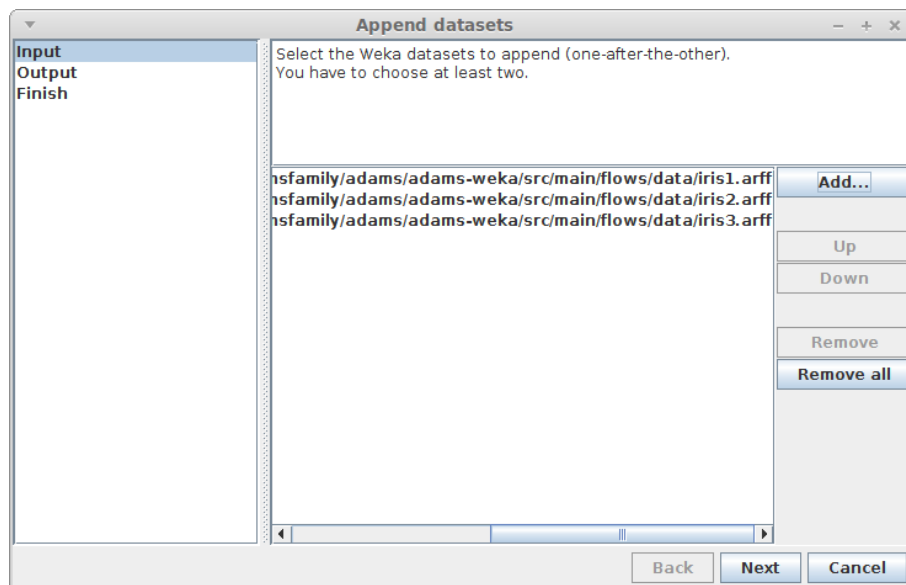



Figure 7.1: Appending datasets wizard.

7.2 Arff viewer

WEKA's *Arff viewer* is available from the ADAMS menu as well.



No.	1: sepalwidth Numeric	2: sepalwidth Numeric	3: petalwidth Numeric	4: petalwidth Numeric	5: class Nominal
1	5.1	3.5	1.4	0.2	Iris-set...
2	4.9	3.0	1.4	0.2	Iris-set...
3	4.7	3.2	1.3	0.2	Iris-set...
4	4.6	3.1	1.5	0.2	Iris-set...
5	5.0	3.6	1.4	0.2	Iris-set...
6	5.4	3.9	1.7	0.4	Iris-set...
7	4.6	3.4	1.4	0.3	Iris-set...
8	5.0	3.4	1.5	0.2	Iris-set...
9	4.4	2.9	1.4	0.2	Iris-set...
10	4.9	3.1	1.5	0.1	Iris-set...
11	5.4	3.7	1.5	0.2	Iris-set...
12	4.8	3.4	1.6	0.2	Iris-set...
13	4.8	3.0	1.4	0.1	Iris-set...
14	4.3	3.0	1.1	0.1	Iris-set...
15	5.8	4.0	1.2	0.2	Iris-set...
16	5.7	4.4	1.5	0.4	Iris-set...
17	5.4	3.9	1.3	0.4	Iris-set...
18	5.1	3.5	1.4	0.3	Iris-set...
19	5.7	3.8	1.7	0.3	Iris-set...
20	5.1	3.8	1.5	0.3	Iris-set...
21	5.4	3.4	1.7	0.2	Iris-set...
22	5.1	3.7	1.5	0.4	Iris-set...
23	4.6	3.6	1.0	0.2	Iris-set...
24	5.1	3.3	1.7	0.5	Iris-set...
25	4.8	3.4	1.9	0.2	Iris-set...
26	5.0	3.0	1.6	0.2	Iris-set...
27	5.0	3.4	1.6	0.4	Iris-set...
28	5.2	3.5	1.5	0.2	Iris-set...
29	5.2	3.4	1.4	0.2	Iris-set...
30	4.7	3.2	1.6	0.2	Iris-set...

Figure 7.2: WEKA's Arff viewer displaying the iris UCI dataset.

7.3 Batch-filter datasets

Training and test set are required to have exactly the same structure in WEKA. If you need to apply a pre-processing filter to both of the files, it is recommended to use *batch-filtering*. This will ensure that any statistics that get calculated based on the first dataset, get applied to second file as well. An example for this is the dictionary of the *StringToWordVector*: the same word dictionary needs to be applied to the second file, in order to generate the same attributes. However, one usually has to resort to the command-line in order to achieve this, which is rather tedious.

The *Batch-filter datasets* wizard allows you to perform the batch-filtering conveniently through the GUI:

- Select the datasets to filter (see 7.3)
- Configure the filter (see 7.4)
- Select the output directory (see 7.5)
- Info dialog after filtering (see 7.6)

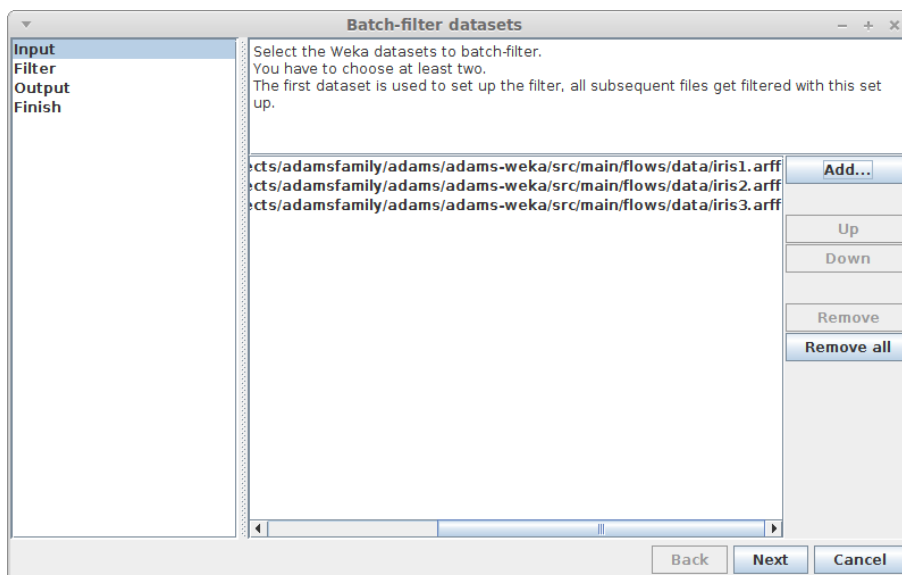


Figure 7.3: Selecting input files for batch-filtering.




Figure 7.4: Configuring the filter.




Figure 7.5: Selecting the output directory for the filtered datasets.




Figure 7.6: The dialog when batch-filtering was successful.

7.4 BayesNet Editor

WEKA's *BayesNet Editor* allows you to view and modify Bayesian nets. Figure 7.7 shows an example net loaded.




Figure 7.7: BayesNet editor displaying a net.

7.5 Dark Lord

The *Dark Lord* wizard allows you to set up and monitor a genetic algorithm that determines the best set of attributes, based on the classifier and metric you choose to optimize on. The output directory that you define (see Figure 7.8), is used to store all the *reduced/optimized* datasets that improved the metric. The optimization can be paused, resumed or stopped with the buttons at the bottom right of the view panel (see Figure 7.9).

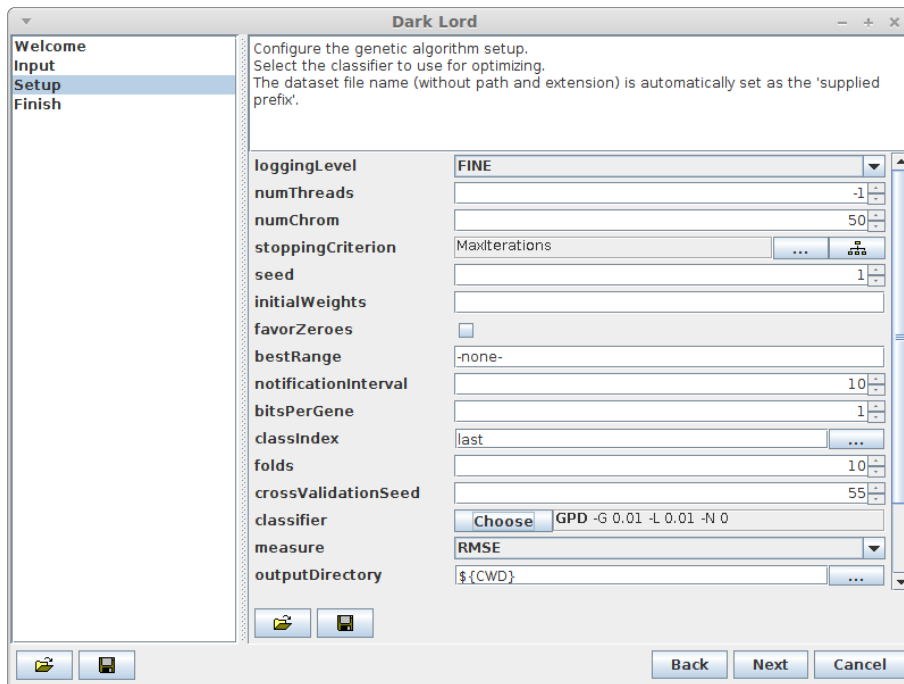


Figure 7.8: Dark Lord setup.




Figure 7.9: Dark Lord run.

7.6 Dataset compatibility

WEKA requires training and test sets to have the same structure, down to the same name and order of nominal labels. Rather than relying on the error message in the Explorer, you can use the *Dataset compatibility* tool to quickly check whether two or more datasets are actually compatible.

The screenshot in Figure 7.10 shows the output when comparing two datasets, one being the original *anneal* UCI dataset and the other one a transformed version.

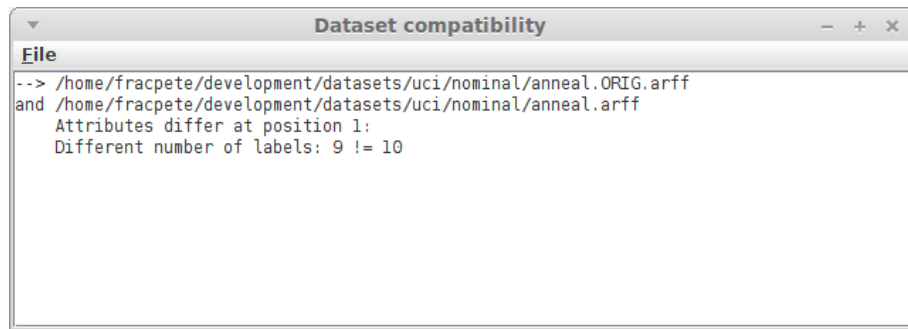


Figure 7.10: Compatibility output for two datasets.

7.7 Explorer

ADAMS contains an extended version of the WEKA Explorer. The interface uses menus instead of buttons to declutter the pre-process tab. Also, it keeps track of the datasets that the user loads, to make re-loading recent files easier. This saves a lot of time when working with the same files on a frequent basis. Furthermore, the user can have an arbitrary number of Explorer sessions in the same window, distinguished by names. Figure 7.11 shows the new interface with the drop-down menu in action.

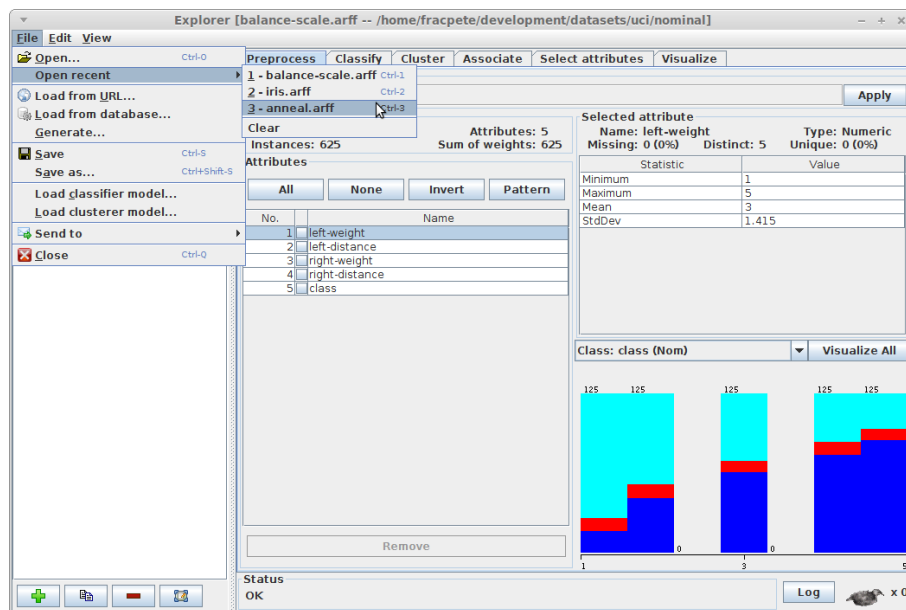


Figure 7.11: Explorer interface with menus.

One very useful feature is the notion of *workspaces* in this interface. You can save the current setup (current dataset, classifiers, clusterers, evaluation set up, results, etc.) to a file and restore all of it in one go again. Unfortunately, not all data can be stored, such as the log, the undo history and the built models or visualizations associated with a results. See Figure 7.12 for the button (highlighted in red) that allows you to load/save workspaces.

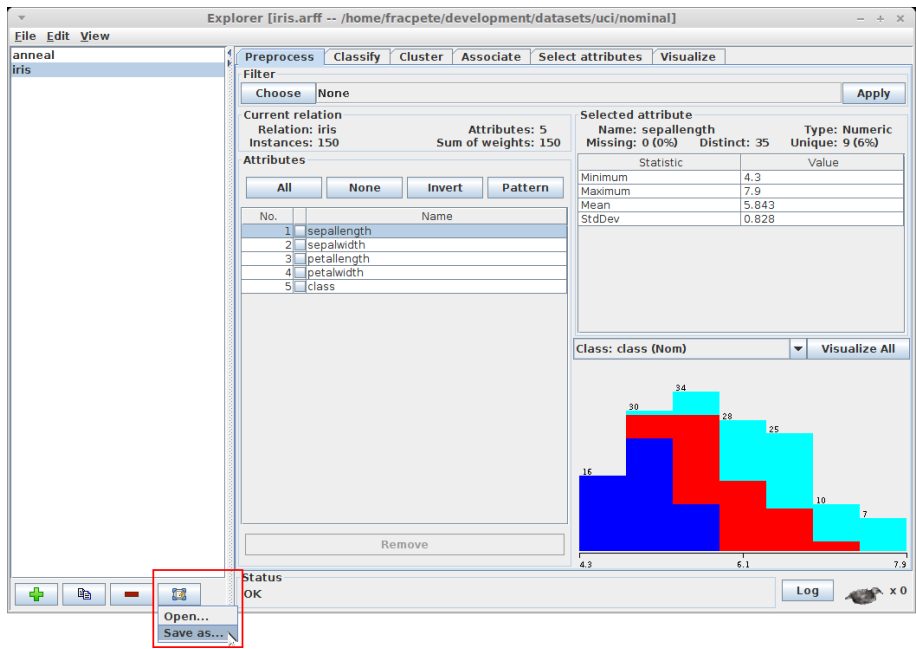


Figure 7.12: Saving/restoring of workspaces.

7.8 Hermione

The *Hermione* wizard allows you to set up and monitor a genetic algorithm that determines the best parameter settings, based on the classifier and metric you choose to optimize on. The output directory that you define (see Figure 7.13), is used to store all the setups/datasets that improved the metric. The optimization can be paused, resumed or stopped with the buttons at the bottom right of the view panel (see Figure 7.14).

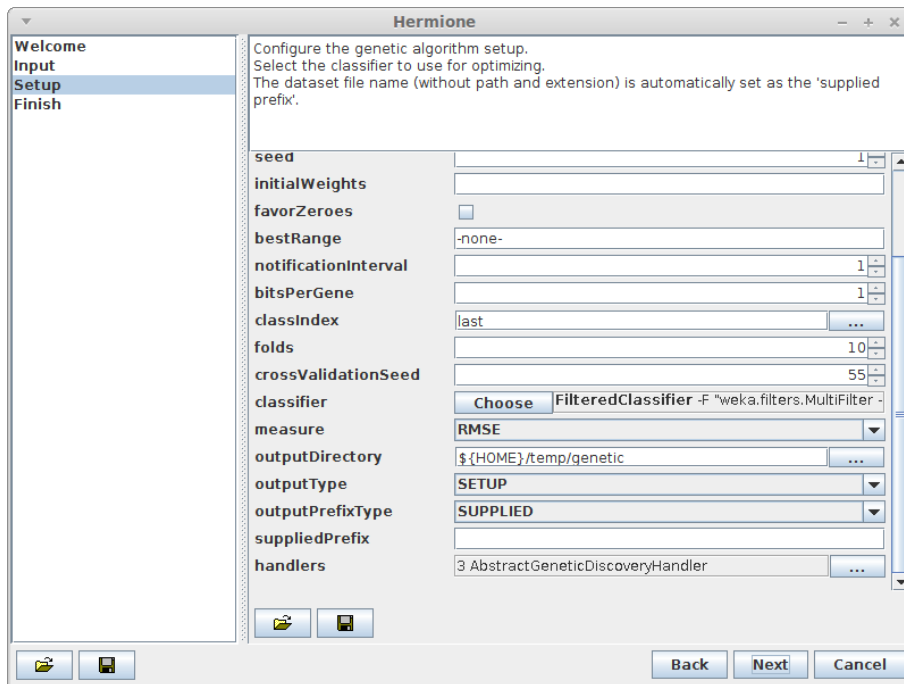


Figure 7.13: Hermione setup.




Figure 7.14: Hermione run.

7.9 Investigator

The *Weka Investigator* (see Figure 7.15) is aimed to be a replacement of the Weka Explorer. The idea is to have a more versatile and efficient user interface: manage multiple sessions, each session can manage multiple datasets (e.g., train and test set), making use of tabs instead of windows (but output tabs can still be detached for comparison), configuration of outputs to generate each time an evaluation is run (statistics, model, classifier errors, etc).




Figure 7.15: Weka Investigator.

7.10 Make compatible datasets

When dealing with data across multiple spreadsheets, e.g., training, test and validation set, then it can happen that these datasets are not compatible in the strict Weka sense.

The *Make compatible datasets* wizard allows you to turn spreadsheets into compatible ARFF files:

- Select the datasets to make compatible (see 7.16)
- Configure how the datasets are being read (see 7.21)
- Select the output directory for the compatible ARFF files (see 7.22)
- Info dialog after processing the datasets (see 7.23)




Figure 7.16: Selecting input files for making compatible datasets.

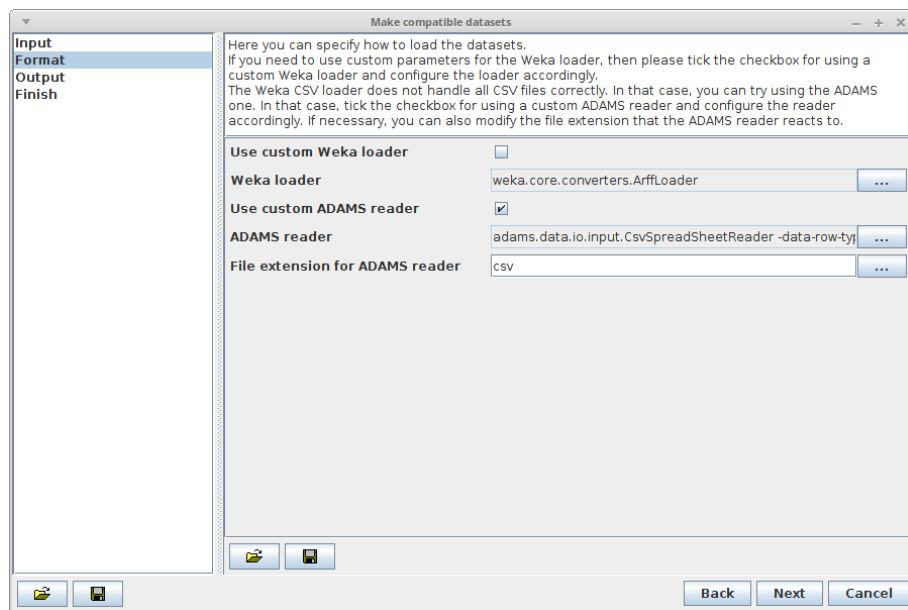


Figure 7.17: Configuring the readers for the datasets.




Figure 7.18: Selecting the output directory for the compatible ARFF files.




Figure 7.19: The dialog when dataset generation was successful.

7.11 Merge datasets

The process of adding reference data to a dataset or performing *data fusion*¹, can be quite often a manual process, copy/pasting datasets in a spreadsheet application.

The *Merge datasets* wizard allows you to perform the process of combining datasets side-by-side conveniently in the GUI:

- Select the datasets to merge (see 7.20)
- Configure the merge (see 7.21)
- Select the output file (see 7.22)
- Info dialog after merging (see 7.23)

Figure 7.24 shows an example of a merged file.

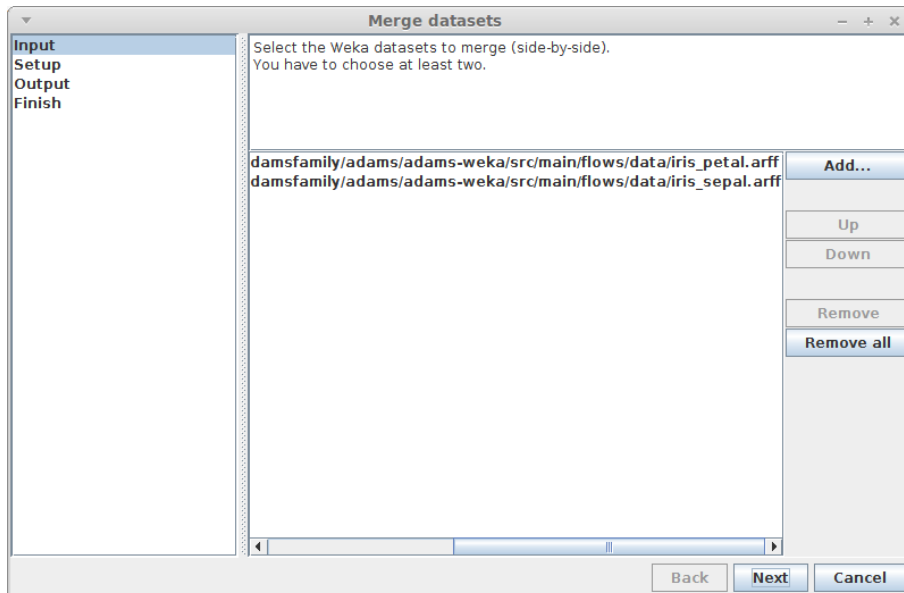


Figure 7.20: Selecting input files for merging.

¹https://en.wikipedia.org/wiki/Data_fusion




Figure 7.21: Configuring the merge.

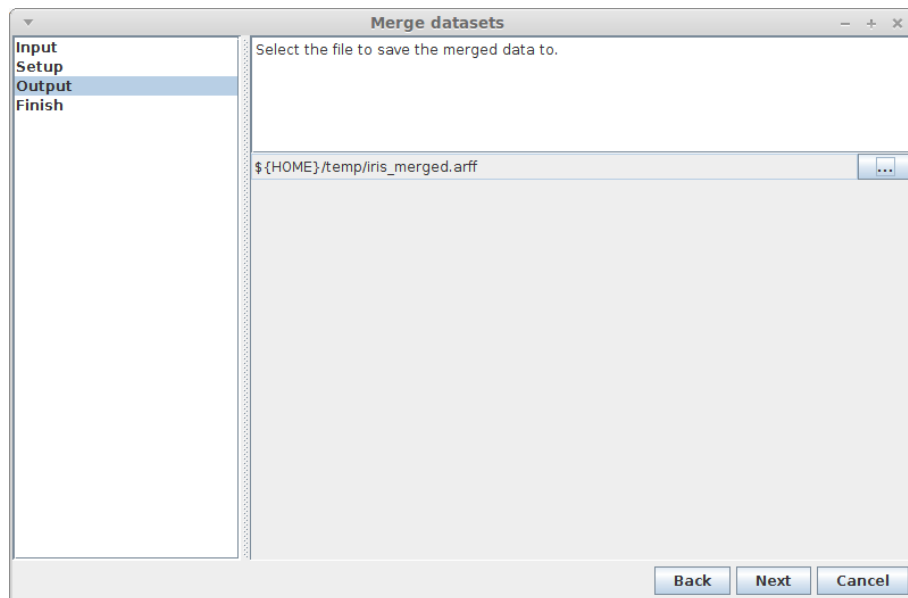


Figure 7.22: Selecting the output directory for the filtered datasets.




Figure 7.23: The dialog when merging was successful.

The screenshot shows the ARFF-Viewer application displaying the merged dataset. The window title is 'ARFF-Viewer - /home/fracpete/temp/iris_merged.arff'. The file 'iris_merged.arff' is loaded. The relation is named 'iris-petal iris-sepal'. The table has 6 columns: No., 1: dataset-1-ID, 2: dataset-1-petalwidth, 3: dataset-1-class, 4: dataset-2-ID, and 5: dataset-2-sep. The data consists of 29 rows of Iris setosa samples.

No.	1: dataset-1-ID	2: dataset-1-petalwidth	3: dataset-1-class	4: dataset-2-ID	5: dataset-2-sep
	Numeric	Numeric	Nominal	Numeric	Numeric
1	1.0	1.4	0.2Iris-setosa	1.0	
2	2.0	1.4	0.2Iris-setosa	2.0	
3	3.0	1.3	0.2Iris-setosa	3.0	
4	4.0	1.5	0.2Iris-setosa	4.0	
5	5.0	1.4	0.2Iris-setosa	5.0	
6	6.0	1.7	0.4Iris-setosa	6.0	
7	7.0	1.4	0.3Iris-setosa	7.0	
8	8.0	1.5	0.2Iris-setosa	8.0	
9	9.0	1.4	0.2Iris-setosa	9.0	
10	10.0	1.5	0.1Iris-setosa	10.0	
11	11.0	1.5	0.2Iris-setosa	11.0	
12	12.0	1.6	0.2Iris-setosa	12.0	
13	13.0	1.4	0.1Iris-setosa	13.0	
14	14.0	1.1	0.1Iris-setosa	14.0	
15	15.0	1.2	0.2Iris-setosa	15.0	
16	16.0	1.5	0.4Iris-setosa	16.0	
17	17.0	1.3	0.4Iris-setosa	17.0	
18	18.0	1.4	0.3Iris-setosa	18.0	
19	19.0	1.7	0.3Iris-setosa	19.0	
20	20.0	1.5	0.3Iris-setosa	20.0	
21	21.0	1.7	0.2Iris-setosa	21.0	
22	22.0	1.5	0.4Iris-setosa	22.0	
23	23.0	1.0	0.2Iris-setosa	23.0	
24	24.0	1.7	0.5Iris-setosa	24.0	
25	25.0	1.9	0.2Iris-setosa	25.0	
26	26.0	1.6	0.2Iris-setosa	26.0	
27	27.0	1.6	0.4Iris-setosa	27.0	
28	28.0	1.5	0.2Iris-setosa	28.0	
29	29.0	1.4	0.2Iris-setosa	29.0	

Figure 7.24: The merged dataset.

7.12 Workbench

The more flexible alternative to the Explorer is called the *Workbench*. A screenshot is shown in Figure 7.25.




Figure 7.25: The Workbench interface.

Bibliography

- [1] *ADAMS* – Advanced Data mining and Machine learning System
<https://adams.cms.waikato.ac.nz/>
- [2] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, Ian H. Witten (2009); *The WEKA Data Mining Software: An Update*; SIGKDD Explorations, Volume 11, Issue 1.
<http://www.cms.waikato.ac.nz/ml/weka/>
- [3] Ian H. Witten, Eibe Frank, Mark A. Hall (2011); *Data Mining: Practical Machine Learning Tools and Techniques*; Third Edition; Morgan Kaufmann; ISBN 978-0-12-374856-0
<http://www.cs.waikato.ac.nz/ml/weka/book.html>
- [4] Barker, G, and McGhie, R (1984) The Biology of Introduced Slugs (Pulmonata) in New Zealand: Introduction and Notes on *Limax Maximus*, NZ Entomologist 8, pp 106-111