arXiv:physics/0703039 [Data Analysis, Statistics and Probability]
CERN-OPEN-2007-007
TMVA version 4.0.1
September 19, 2009
http:// tmva.sourceforge.net

TMVA 4

Toolkit for Multivariate Data Analysis with ROOT

Users Guide

A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, H. Voss

Contributed to TMVA have:

M. Backes, T. Carli, O. Cohen, A. Christov, D. Dannheim, K. Danielowski,
S. Henrot-Versillé, M. Jachowski, K. Kraszewski, A. Krasznahorkay Jr.,
M. Kruk, Y. Mahalalel, R. Ospanov, X. Prudent, A. Robert, D. Schouten,
F. Tegenfeldt, A. Voigt, K. Voss, M. Wolter, A. Zemla

Abstract — In high-energy physics, with the search for ever smaller signals in ever larger data sets, it has become essential to extract a maximum of the available information from the data. Multivariate classification methods based on machine learning techniques have become a fundamental ingredient to most analyses. Also the multivariate classifiers themselves have significantly evolved in recent years. Statisticians have found new ways to tune and to combine classifiers to further gain in performance. Integrated into the analysis framework ROOT, TMVA is a toolkit which hosts a large variety of multivariate classification algorithms. Training, testing, performance evaluation and application of all available classifiers is carried out simultaneously via user-friendly interfaces. With version 4, TMVA has been extended to multivariate regression of a real-valued target vector. Regression is invoked through the same user interfaces as classification. TMVA 4 also features more flexible data handling allowing one to arbitrarily form combined MVA methods. A generalised boosting method is the first realisation benefiting from the new framework.

TMVA 4.0.1 – Toolkit for Multivariate Data Analysis with ROOT

Copyright © 2005-2009, Regents of CERN (Switzerland), DESY (Germany), MPI-Kernphysik Heidelberg (Germany), University of Bonn (Germany), and University of Victoria (Canada).
BSD license: http://tmva.sourceforge.net/LICENSE.

Authors:

Andreas Hoecker (CERN, Switzerland) ⟨andreas.hoecker@cern.ch⟩,
Peter Speckmayer (CERN, Switzerland) ⟨peter.speckmayer@cern.ch⟩,
Jörg Stelzer (CERN, Switzerland) ⟨joerg.stelzer@cern.ch⟩,
Jan Therhaag (Universität Bonn, Germany) ⟨therhaag@physik.uni-bonn.de⟩,
Eckhard von Toerne (Universität Bonn, Germany) ⟨evt@physik.uni-bonn.de⟩,
Helge Voss (MPI für Kernphysik Heidelberg, Germany) ⟨helge.voss@cern.ch⟩,
Moritz Backes (Geneva University, Switzerland) moritz.backes@cern.ch,
Tancredi Carli (CERN, Switzerland) ⟨tancredi.carli@cern.ch⟩,
Or Cohen (CERN, Switzerland and Technion, Israel) ⟨or.cohen@cern.ch⟩,
Asen Christov (Universität Freiburg, Germany) ⟨christov@physik.uni-freiburg.de⟩,
Krzysztof Danielowski (IFJ and AGH/UJ, Krakow, Poland) ⟨Krzysztof.Danielowski@cern.ch⟩,
Dominik Dannheim (CERN, Switzerland) ⟨Dominik.Dannheim@cern.ch⟩,
Sophie Henrot-Versillé (LAL Orsay, France) ⟨versille@lal.in2p3.fr⟩,
Matthew Jachowski (Stanford University, USA) ⟨jachowski@stanford.edu⟩,
Kamil Kraszewski (IFJ and AGH/UJ, Krakow, Poland) ⟨kamil.bartlomiej.kraszewski@cern.ch⟩,
Attila Krasznahorkay Jr. (CERN, CH, and Manchester U., UK) ⟨Attila.Krasznahorkay@cern.ch⟩,
Maciej Kruk (IFJ and AGH/UJ, Krakow, Poland) ⟨maciej.mateusz.kruk@cern.ch⟩,
Yair Mahalalel (Tel Aviv University, Israel) ⟨yair@mahalalel.com⟩,
Rustem Ospanov (University of Texas, USA) ⟨rustem@fnal.gov⟩,
Xavier Prudent (LAPP Annecy, France) ⟨prudent@lapp.in2p3.fr⟩,
Doug Schouten (S. Fraser U., Canada) ⟨dschoute@sfu.ca⟩,
Fredrik Tegenfeldt (Iowa University, USA) ⟨fredrik.tegenfeldt@cern.ch⟩,
Arnaud Robert (LPNHE Paris, France) ⟨arobert@lpnhe.in2p3.fr⟩,
Alexander Voigt (CERN, Switzerland) ⟨alexander.voigt@cern.ch⟩,
Kai Voss (University of Victoria, Canada) ⟨kai.voss@cern.ch⟩,
Marcin Wolter (IFJ PAN Krakow, Poland) ⟨marcin.wolter@ifj.edu.pl⟩,
Andrzej Zemla (IFJ PAN Krakow, Poland) ⟨zemla@corcoran.if.uj.edu.pl⟩,
and valuable contributions from many users, please see acknowledgements.

Contents

1 Introduction **1**

Copyrights and credits 3

2 TMVA Quick Start **4**

2.1 How to download and build TMVA 4

2.2 Version compatibility 5

2.3 Avoiding conflicts between external TMVA and ROOT's internal one 5

2.4 The TMVA namespace 5

2.5 Example jobs 5

2.6 Running the example 6

2.7 Displaying the results 7

2.8 Getting help 10

3 Using TMVA **12**

3.1 The TMVA Factory 13

3.1.1 Specifying training and test data . . . 15

3.1.2 Defining input variables, targets and event weights 17

3.1.3 Negative event weights 19

3.1.4 Preparing the training and test data . 19

3.1.5 Booking MVA methods 21

3.1.6 Help option for MVA booking 22

3.1.7 Training the MVA methods 22

3.1.8 Testing the MVA methods 23

3.1.9 Evaluating the MVA methods 23

3.1.10 Classification performance evaluation 23

3.1.11 Regression performance evaluation . 24

3.1.12 Overtraining 27

3.1.13 Other representations of MVA outputs for classification: probabilities and *Rarity* 28

3.2 ROOT macros to plot training, testing and evaluation results 29

3.3 The TMVA Reader 31

3.3.1 Specifying input variables 31

3.3.2 Booking MVA methods 33

3.3.3 Requesting the MVA response 33

3.4 An alternative to the Reader: standalone C++ response classes 34

4 Data Preprocessing **37**

4.1 Transforming input variables 37

4.1.1 Variable normalisation 38

4.1.2 Variable decorrelation 38

4.1.3 Principal component decomposition . 39

4.1.4 Gaussian transformation of variables ("Gaussianisation") 40

4.1.5 Booking and chaining transformations 40

4.2 Binary search trees 41

5 Probability Density Functions – the *PDF* Class **41**

5.1 Nonparametric PDF fitting using spline functions 42

5.2 Nonparametric PDF parameterisation using kernel density estimators 44

6 Optimisation and Fitting **45**

6.1 Monte Carlo sampling 46

6.2 Minuit minimisation 46

6.3 Genetic Algorithm 47

6.4 Simulated Annealing 50

6.5 Combined fitters 51

7 Boosting and Bagging **52**

7.1 Adaptive Boost (AdaBoost) 52

7.2 Gradient Boost 54

7.3 Bagging 55

8 The TMVA Methods **57**

8.1 Rectangular cut optimisation 57

8.1.1 Booking options 59

8.1.2 Description and implementation . . . 60

8.1.3 Variable ranking 61

8.1.4 Performance 61

8.2 Projective likelihood estimator (PDE approach) 62

8.2.1 Booking options 62

8.2.2 Description and implementation . . . 62

8.2.3 Variable ranking 63

8.2.4 Performance 64

8.3 Multidimensional likelihood estimator (PDE range-search approach) 64

8.3.1 Booking options 65

8.3.2 Description and implementation . . . 65

8.3.3 Variable ranking 68

8.3.4 Performance 69

8.4 Likelihood estimator using self-adapting phase-space binning (PDE-Foam) 69

8.4.1 Booking options 69

8.4.2 Description and implementation of the foam algorithm . . . 70
8.4.3 Classification . . . 74
8.4.4 Regression . . . 76
8.4.5 Visualisation of the foam via projections to 2 dimensions . . . 78
8.4.6 Performance . . . 79
8.5 k-Nearest Neighbour (k-NN) Classifier . . . 79
8.5.1 Booking options . . . 79
8.5.2 Description and implementation . . . 80
8.5.3 Ranking . . . 82
8.5.4 Performance . . . 82
8.6 H-Matrix discriminant . . . 83
8.6.1 Booking options . . . 83
8.6.2 Description and implementation . . . 83
8.6.3 Variable ranking . . . 84
8.6.4 Performance . . . 84
8.7 Fisher discriminants (linear discriminant analysis) . . . 84
8.7.1 Booking options . . . 84
8.7.2 Description and implementation . . . 85
8.7.3 Variable ranking . . . 86
8.7.4 Performance . . . 86
8.8 Linear discriminant analysis (LD) . . . 86
8.8.1 Booking options . . . 86
8.8.2 Description and implementation . . . 87
8.8.3 Variable ranking . . . 88
8.8.4 Regression with LD . . . 88
8.8.5 Performance . . . 88
8.9 Function discriminant analysis (FDA) . . . 88
8.9.1 Booking options . . . 89
8.9.2 Description and implementation . . . 90
8.9.3 Variable ranking . . . 90
8.9.4 Performance . . . 90
8.10 Artificial Neural Networks (nonlinear discriminant analysis) . . . 91
8.10.1 Booking options . . . 91
8.10.2 Description and implementation . . . 95
8.10.3 Network architecture . . . 96
8.10.4 Training of the neural network . . . 97
8.10.5 Variable ranking . . . 99
8.10.6 Performance . . . 99
8.11 Support Vector Machine (SVM) . . . 99
8.11.1 Booking options . . . 100
8.11.2 Description and implementation . . . 101
8.11.3 Variable ranking . . . 104
8.11.4 Performance . . . 104
8.12 Boosted Decision and Regression Trees . . . 104
8.12.1 Booking options . . . 104
8.12.2 Description and implementation . . . 105
8.12.3 Boosting, Bagging and Randomising 108
8.12.4 Variable ranking . . . 110
8.12.5 Performance . . . 111
8.13 Predictive learning via rule ensembles (RuleFit) . . . 111
8.13.1 Booking options . . . 112
8.13.2 Description and implementation . . . 112
8.13.3 Variable ranking . . . 115
8.13.4 Friedman's module . . . 117
8.13.5 Performance . . . 117
9 Combining MVA Methods 118
9.1 Boosted classifiers . . . 119
9.1.1 Booking options . . . 119
9.1.2 Boostable classifiers . . . 120
9.1.3 Monitoring tools . . . 121
9.1.4 Variable ranking . . . 121
9.2 Category Classifier . . . 121
9.2.1 Booking options . . . 122
9.2.2 Description and implementation . . . 123
9.2.3 Variable ranking . . . 124
9.2.4 Performance . . . 124
10 Which MVA method should I use for my problem? 126
11 TMVA implementation status summary for classification and regression 127
12 Conclusions and Plans 129
Acknowledgements 133
A More Classifier Booking Examples 134
Bibliography 137
Index 139

1 Introduction

The Toolkit for Multivariate Analysis (TMVA) provides a ROOT-integrated [1] environment for the processing, parallel evaluation and application of multivariate classification and – since TMVA version 4 – multivariate regression techniques.[1] All multivariate techniques in TMVA belong to the family of "supervised learnning" algorithms. They make use of training events, for which the desired output is known, to determine the mapping function that either discribes a decision boundary (classification) or an approximation of the underlying functional behaviour defining the target value (regression). The mapping function can contain various degrees of approximations and may be a single global function, or a set of local models. TMVA is specifically designed for the needs of high-energy physics (HEP) applications, but should not be restricted to these. The package includes:

- Rectangular cut optimisation (binary splits, Sec. 8.1).
- Projective likelihood estimation (Sec. 8.2).
- Multi-dimensional likelihood estimation (PDE range-search – Sec. 8.3, PDE-Foam – Sec. 8.4, and k-NN – Sec. 8.5).
- Linear and nonlinear discriminant analysis (H-Matrix – Sec. 8.6, Fisher – Sec. 8.7, LD – Sec. 8.8, FDA – Sec. 8.9).
- Artificial neural networks (three different multilayer perceptron implementations – Sec. 8.10).
- Support vector machine (Sec. 8.11).
- Boosted/bagged decision trees (Sec. 8.12).
- Predictive learning via rule ensembles (RuleFit, Sec. 8.13).
- A generic boost classifier allowing one to boost any of the above classifiers (Sec. 9).
- A generic category classifier allowing one to split the training data into disjoint categories with independent MVAs.

The software package consists of abstract, object-oriented implementations in C++/ROOT for each of these multivariate analysis (MVA) techniques, as well as auxiliary tools such as parameter fitting and transformations. It provides training, testing and performance evaluation algorithms

[1] A classification problem corresponds in more general terms to a *discretised regression* problem. A regression is the process that estimates the parameter values of a function, which predicts the value of a response variable (or vector) in terms of the values of other variables (the *input* variables). A typical regression problem in High-Energy Physics is for example the estimation of the energy of a (hadronic) calorimeter cluster from the cluster's electromagnetic cell energies. The user provides a single dataset that contains the input variables and one or more target variables. The interface to defining the input and target variables, the booking of the multivariate methods, their training and testing is very similar to the syntax in classification problems. Communication between the user and TMVA proceeds conveniently via the Factory and Reader classes. Due to their similarity, classification and regression are introduced together in this Users Guide. Where necessary, differences are pointed out.

and visualisation scripts. Detailed descriptions of all the TMVA methods and their options for classification and (where available) regression tasks are given in Sec. 8. Their training and testing is performed with the use of user-supplied data sets in form of ROOT trees or text files, where each event can have an individual weight. The true sample composition (for event classification) or target value (for regression) in these data sets must be supplied for each event. Preselection requirements and transformations can be applied to input data. TMVA supports the use of variable combinations and formulas with a functionality similar to the one available for the `Draw` command of a ROOT tree.

TMVA works in transparent factory mode to guarantee an unbiased performance comparison between MVA methods: they all see the same training and test data, and are evaluated following the same prescriptions within the same execution job. A *Factory* class organises the interaction between the user and the TMVA analysis steps. It performs preanalysis and preprocessing of the training data to assess basic properties of the discriminating variables used as inputs to the classifiers. The linear correlation coefficients of the input variables are calculated and displayed. For regression, also nonlinear correlation measures are given, such as the correlation ratio and mutual information between input variables and output target. A preliminary ranking is derived, which is later superseded by algorithm-specific variable rankings. For classification problems, the variables can be linearly transformed (individually for each MVA method) into a non-correlated variable space, projected upon their principle components, or transformed into a normalised Gaussian shape. Transformations can also be arbitrarily concatenated.

To compare the signal-efficiency and background-rejection performance of the classifiers, or the average variance between regression target and estimation, the analysis job prints – among other criteria – tabulated results for some benchmark values (see Sec. 3.1.9). Moreover, a variety of graphical evaluation information acquired during the training, testing and evaluation phases is stored in a ROOT output file. These results can be displayed using macros, which are conveniently executed via graphical user interfaces (each one for classification and regression) that come with the TMVA distribution (see Sec. 3.2).

The TMVA training job runs alternatively as a ROOT script, as a standalone executable, or as a python script via the PyROOT interface. Each MVA method trained in one of these applications writes its configuration and training results in a result ("weight") file, which in the default configuration has human readable XML format.

A light-weight *Reader* class is provided, which reads and interprets the weight files (interfaced by the corresponding method), and which can be included in any C++ executable, ROOT macro, or python analysis job (see Sec. 3.3).

For standalone use of the trained MVA method, TMVA also generates lightweight C++ response classes (not available for all methods), which contain the encoded information from the weight files so that these are not required anymore. These classes do not depend on TMVA or ROOT, neither on any other external library (see Sec. 3.4).

We have put emphasis on the clarity and functionality of the Factory and Reader interfaces to the user applications, which will hardly exceed a few lines of code. All MVA methods run with reasonable default configurations and should have satisfying performance for average applications. *We stress*

however that, to solve a concrete problem, all methods require at least some specific tuning to deploy their maximum classification or regression capabilities. Individual optimisation and customisation of the classifiers is achieved via configuration strings when booking a method.

This manual introduces the TMVA Factory and Reader interfaces, and describes design and implementation of the MVA methods. It is not the aim here to provide a general introduction to MVA techniques. Other excellent reviews exist on this subject (see, e.g., Refs. [2–4]). The document begins with a quick TMVA start reference in Sec. 2, and provides a more complete introduction to the TMVA design and its functionality for both, classification and regression analyses in Sec. 3. Data preprocessing such as the transformation of input variables and event sorting are discussed in Sec. 4. In Sec. 5, we describe the techniques used to estimate probability density functions from the training data. Section 6 introduces optimisation and fitting tools commonly used by the methods. All the TMVA methods including their configurations and tuning options are described in Secs. 8.1–8.13. Guidance on which MVA method to use for varying problems and input conditions is given in Sec. 10. An overall summary of the implementation status of all TMVA methods is provided in Sec. 11.

Copyrights and credits

TMVA is an open source product. Redistribution and use of TMVA in source and binary forms, with or without modification, are permitted according to the terms listed in the BSD license.[2] Several similar combined multivariate analysis ("machine learning") packages exist with rising importance in most fields of science and industry. In the HEP community the package *StatPatternRecognition* [5, 6] is in use (for classification problems only). The idea of parallel training and evaluation of MVA-based classification in HEP has been pioneered by the *Cornelius* package, developed by the Tagging Group of the BABAR Collaboration [7]. See further credits and acknowledgments on page 133.

[2]For the BSD l icense, see http://tmva.sourceforge.net/LICENSE.

2 TMVA Quick Start

To run TMVA it is not necessary to know much about its concepts or to understand the detailed functionality of the multivariate methods. Better, just begin with the quick start tutorial given below. One should note that the TMVA version obtained from the open source software platform Sourceforge.net (where TMVA is hosted), and the one included in ROOT, have different directory structures for the example macros used for the tutorial. Wherever differences in command lines occur, they are given for both versions.

Classification and regression analyses in TMVA have similar training, testing and evaluation phases, and will be treated in parallel in the following.

2.1 How to download and build TMVA

TMVA is developed and maintained at Sourceforge.net (http://tmva.sourceforge.net). It is built upon ROOT (http://root.cern.ch/), so that for TMVA to run ROOT must be installed. Since ROOT version 5.11/06, TMVA comes as integral part of ROOT and can be used from the ROOT prompt without further preparation. For older ROOT versions or *if the latest TMVA features are desired*, the TMVA source code can be downloaded from Sourceforge.net. Since we do not provide prebuilt libraries for any platform, the library must be built by the user (easy – see below). The source code can be either downloaded as a gzipped tar file or via (anonymous) SVN access:

```
~> svn co https://tmva.svn.sourceforge.net/svnroot/tmva/tags/V04-00-01/TMVA \
         TMVA-4.0.1
```

Code Example 1: Source code download via SVN. The latest version (SVN trunk) can be downloaded by typing the same command without specifying a version: `svn co http:://...tmva/trunk/TMVA`. For the latest TMVA version see http://tmva.sourceforge.net/.

While the source code is known to compile with VisualC++ on Windows (which is a requirement for ROOT), we do not provide project support for this platform yet. For Unix and most Linux flavours custom Makefiles are provided with the TMVA distribution, so that the library can be built by typing:

```
~> cd TMVA
~/TMVA> source setup.sh # for c-shell family: source setup.csh
~/TMVA> cd src
~/TMVA/src> make
```

Code Example 2: Building the TMVA library under Linux/Unix using the provided Makefile. The `setup.[c]sh` script must be executed to ensure the correct setting of symbolic links and library paths required by TMVA.

After compilation, the library `TMVA/lib/libTMVA.1.so` should be present.

2.2 Version compatibility

TMVA can be run with any ROOT version equal or above v5.08. The few occurring conflicts due to ROOT source code evolution after v5.08 are intercepted in TMVA via C++ preprocessor conditions.

2.3 Avoiding conflicts between external TMVA and ROOT's internal one

To use a more recent version of TMVA than the one present in the local ROOT installation, one needs to download the desired TMVA release from Sourceforge.net, to compile it against the local ROOT version, and to make sure the newly built library `TMVA/lib/libTMVA.1.so` is used instead of ROOT's internal one. When running TMVA in a CINT macro the new library must be loaded first via: `gSystem->Load("TMVA/lib/libTMVA.1")`. This can be done directly in the macro or in a file that is automatically loaded at the start of CINT (for an example, see the files `.rootrc` and `TMVAlogon.C` in the `TMVA/macros/` directory). When running TMVA in an executable, the corresponding shared library needs to be linked. Once this is done, ROOT's own `libTMVA.so` library will not be invoked anymore.

2.4 The TMVA namespace

All TMVA classes are embedded in the namespace `TMVA`. For interactive access, or use in macros the classes must thus be preceded by `TMVA::`, or one may use the command `using namespace TMVA` instead.

2.5 Example jobs

TMVA comes with example jobs for the training phase (this phase actually includes training, testing and evaluation) using the TMVA *Factory*, as well as the application of the training results in a classification or regression analysis using the TMVA *Reader*. The first task is performed in the programs `TMVAClassification` or `TMVARegression`, respectively, and the second task in `TMVAClassificationApplication` or `TMVARegressionApplication`.

In the ROOT version of TMVA the above macros (extension `.C`) are located in the directory `$ROOTSYS/tmva/test/`.

In the Sourceforge.net version these macros are located in `TMVA/macros/`. At Sourceforge.net we also provide these examples in form of the C++ executables (replace `.C` by `.cxx`), which are located in `TMVA/execs/`. To build the executables, type `cd ~/TMVA/execs/; make`, and then simply execute them by typing `./TMVAClassification` or `./TMVARegression` (and similarly for the applications). To illustrate how TMVA can be used in a python script via PyROOT we also provide the script

TMVAClassification.py located in TMVA/python/, which has the same functionality as the macro TMVAClassification.C (the other macros are not provided as python scripts).

2.6 Running the example

The most straightforward way to get started with TMVA is to simply run the TMVAClassification.C or TMVARegression.C example macros. Both use academic toy datasets for training and testing, which, for classification, consists of four linearly correlated, Gaussian distributed discriminating input variables, with different sample means for signal and background, and, for regression, has two input variables with fuzzy parabolic dependence on the target (fvalue), and no correlations among themselves. All classifiers are trained, tested and evaluated using the toy datasets in the same way the user is expected to proceed for his or her own data. It is a valuable exercise to look at the example file in more detail. Most of the command lines therein should be self explaining, and one will easily find how they need to be customized to apply TMVA to a real use case. A detailed description is given in Sec. 3.

The toy datasets used by the examples are included in the Sourceforge.net download. For the ROOT distribution, the macros automatically fetch the data file from the web using the corresponding TFile constructor, e.g., TFile::Open("http://root.cern.ch/files/tmva_class_example.root") for classification (tmva_reg_example.root for regression). The example ROOT macros can be run directly in the TMVA/macros/ directory (Sourceforge.net), or in any designated test directory workdir, after adding the macro directory to ROOT's macro search path:

```
~/workdir> echo "Unix.*.Root.MacroPath: ~/TMVA/macros" >> .rootrc
~/workdir> root -l ~/TMVA/macros/TMVAClassification.C
```

Code Example 3: Running the example TMVAClassification.C using the Sourceforge.net version of TMVA (similarly for TMVARegression.C).

```
~/workdir> echo "Unix.*.Root.MacroPath: $ROOTSYS/tmva/test" >> .rootrc
~/workdir> root -l $ROOTSYS/tmva/test/TMVAClassification.C
```

Code Example 4: Running the example TMVAClassification.C using the ROOT version of TMVA (similarly for TMVARegression.C).

It is also possible to explicitly select the MVA methods to be processed (here an example given for a classification task with the Sourceforge.net version):

```
~/workdir> root -l ~/TMVA/macros/TMVAClassification.C\(\"Fisher,Likelihood\"\)
```

Code Example 5: Running the example `TMVAClassification.C` and processing only the Fisher and likelihood classifiers. Note that the backslashes are mandatory. The macro `TMVARegression.C` can be called accordingly.

where the names of the MVA methods are predifined in the macro.

The training job provides formatted output logging containing analysis information such as: linear correlation matrices for the input variables, correlation ratios and mutual information (see below) between input variables and regression targets, variable ranking, summaries of the MVA configurations, goodness-of-fit evaluation for PDFs (if requested), signal and background (or regression target) correlations between the various MVA methods, decision overlaps, signal efficiencies at benchmark background rejection rates (classification) or deviations from target (regression), as well as other performance estimators. Comparison between the results for training and independent test samples provides overtraining validation.

2.7 Displaying the results

Besides so-called "weight" files containing the method-specific training results, TMVA also provides a variety of control and performance plots that can be displayed via a set of ROOT macros available in `TMVA/macros/` or `$ROOTSYS/tmva/test/` for the Sourceforge.net and ROOT distributions of TMVA, respectively. The macros are summarized in Tables 2 and 4 on page 30. At the end of the example jobs a graphical user interface (GUI) is displayed, which conveniently allows to run these macros (see Fig. 1).

Examples for plots produced by these macros are given in Figs. 3–5 for a classification problem. The distributions of the input variables for signal and background according to our example job are shown in Fig. 2. It is useful to quantify the correlations between the input variables. These are drawn in form of a scatter plot with the superimposed profile for two of the input variables in Fig. 3 (upper left). As will be discussed in Sec. 4, TMVA allows to perform a linear decorrelation transformation of the input variables prior to the MVA training (for classification only). The result of such decorrelation is shown at the upper right hand plot of Fig. 3. The lower plots display the linear correlation coefficients between all input variables, for the signal and background training samples of the classification example.

Figure 4 shows several classifier output distributions for signal and background events based on the test sample. By TMVA convention, signal (background) events accumulate at large (small) classifier output values. Hence, cutting on the output and retaining the events with y larger than the cut requirement selects signal samples with efficiencies and purities that respectively decrease and increase with the cut value. The resulting relations between background rejection versus signal efficiency are shown in Fig. 5 for all classifiers that were used in the example macro. This plot belongs to the class of *Receiver Operating Characteristic* (ROC) diagrams, which in its standard

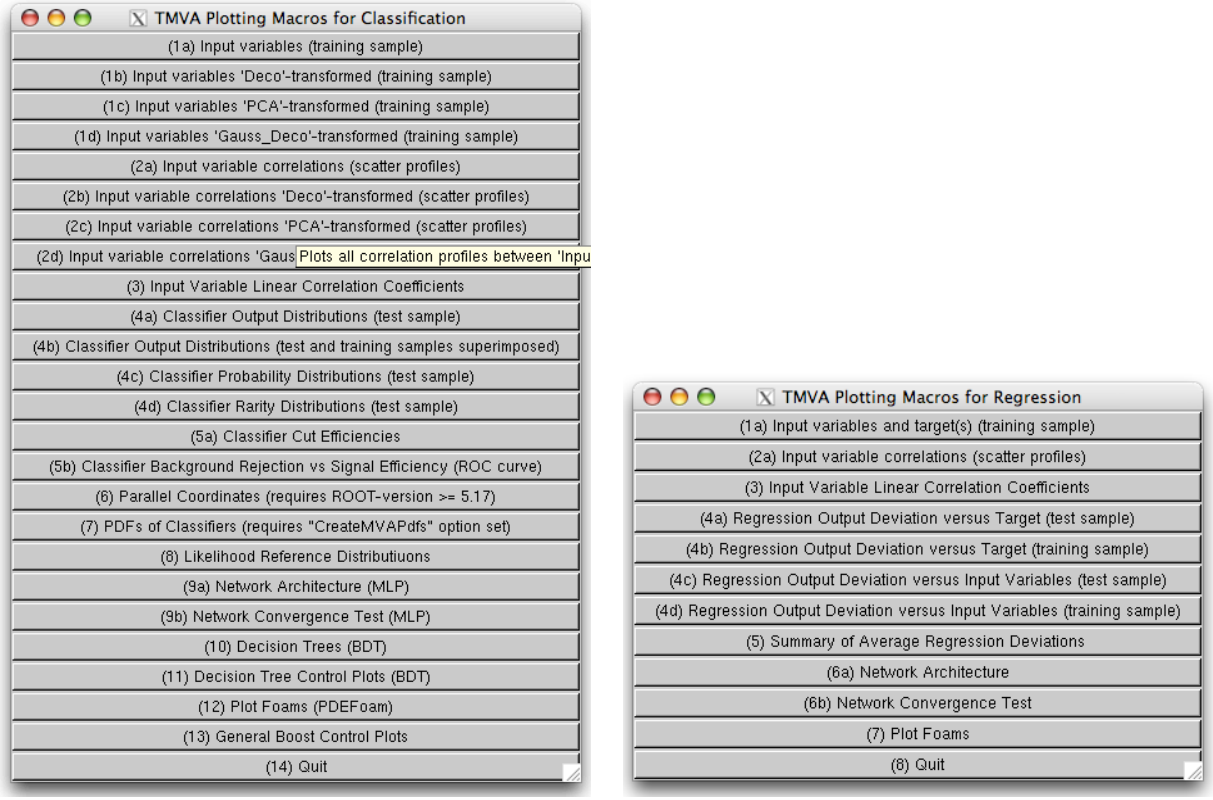

Figure 1: Graphical user interfaces (GUI) to execute macros displaying training, test and evaluation results (cf. Tables 2 and 4 on page 30) for classification (left) and regression problems (right). The classification GUI can be launched manually by executing the scripts `TMVA/macros/TMVAGui.C` (Sourceforge.net version) or `$ROOTSYS/tmva/test/TMVAGui.C` (ROOT version) in a ROOT session. To launch the regression GUI use the macro `TMVARegGui.C`.

Classification (left). The buttons behave as follows: (1a) plots the signal and background distributions of input variables (training sample), (1b–d) the same after applying the corresponding preprocessing transformation of the input variables, (2a–f) scatter plots with superimposed profiles for all pairs of input variables for signal and background and the applied transformations (training sample), (3) correlation coefficients between the input variables for signal and background (training sample), (4a/b) signal and background distributions for the trained classifiers (test sample/test and training samples superimposed to probe overtraining), (4c,d) the corresponding probability and Rarity distributions of the classifiers (where requested, cf. see Sec. 3.1.13), (5a) signal and background efficiencies and purities versus the cut on the classifier output for the expected numbers of signal and background events (before applying the cut) given by the user (an input dialog box pops up, where the numbers are inserted), (5b) background rejection versus signal efficiency obtained when cutting on the classifier outputs (ROC curve, from the test sample), (6) plot of so-called Parallel Coordinates visualising the correlations among the input variables, and among the classifier and the input variables, (7–13) show classifier specific diagnostic plots, and (14) quits the GUI. Titles greyed out indicate actions that are not available because the corresponding classifier has not been trained or because the transformation was not requested.

Regression (right). The buttons behave as follows: (1–3) same as for classification GUI, (4a–d) show the linear deviations between regression targets and estimates versus the targets or input variables for the test and training samples, respectively, (5) compares the average deviations between target and MVA output for the trained methods, and (6–8) are as for the classification GUI.

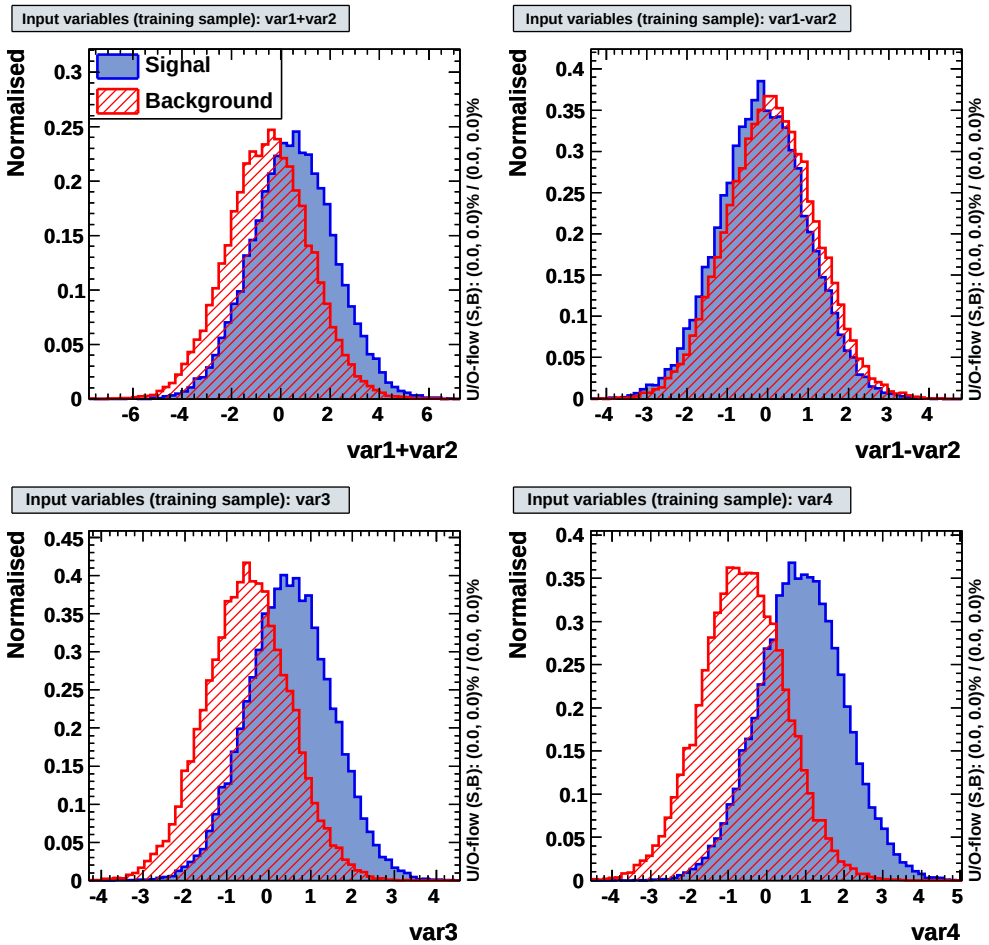

Figure 2: Example plots for input variable distributions. The histogram limits are chosen to zoom into the bulk of the distributions, which may lead to truncated tails. The vertical text on the right-hand side of the plots indicates the under- and overflows. The limits in terms of multiples of the distribution's RMS can be adjusted in the user script by modifying the variable `(TMVA::gConfig().GetVariablePlotting()).fTimesRMS` (cf. Code Example 20).

form shows the true positive rate versus the false positive rate for the different possible cutpoints of a hypothesis test.

As an example for multivariate regression, Fig. 6 displays the deviation between the regression output and target values for linear and nonlinear regression algorithms.

More macros are available to validate training and response of specific MVA methods. For example, the macro `likelihoodrefs.C` compares the probability density functions used by the likelihood classifier to the normalised variable distributions of the training sample. It is also possible to visualize the MLP neural network architecture and to draw decision trees (see Table 4).

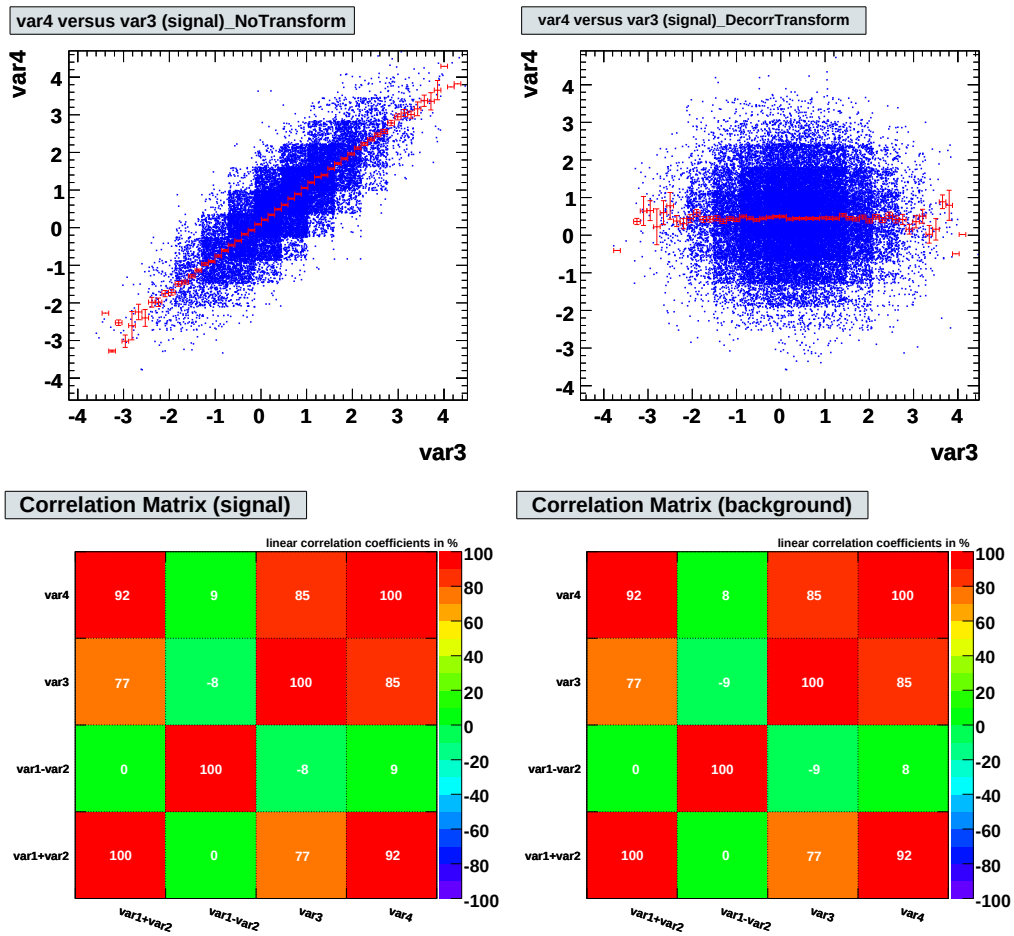

Figure 3: Correlation between input variables. Upper left: correlations between var3 and var4 for the signal training sample. Upper right: the same after applying a linear decorrelation transformation (see Sec. 4.1.2). Lower plots: linear correlation coefficients for the signal and background training samples.

2.8 Getting help

Several help sources exist for TMVA (all web address given below are also linked from the TMVA home page http://tmva.sourceforge.net).

- Information on how to download and install TMVA, and the TMVA *Quick-start* commands are also available on the web at: http://tmva.sourceforge.net/howto.shtml.
- TMVA tutorial: https://twiki.cern.ch/twiki/bin/view/TMVA.
- An up-to-date reference of all configuration options for the TMVA Factory, the fitters, and all the MVA methods: http://tmva.sourceforge.net/optionRef.html.
- On request, the TMVA methods provide a help message with a brief description of the method, and hints for improving the performance by tuning the available configuration options. The

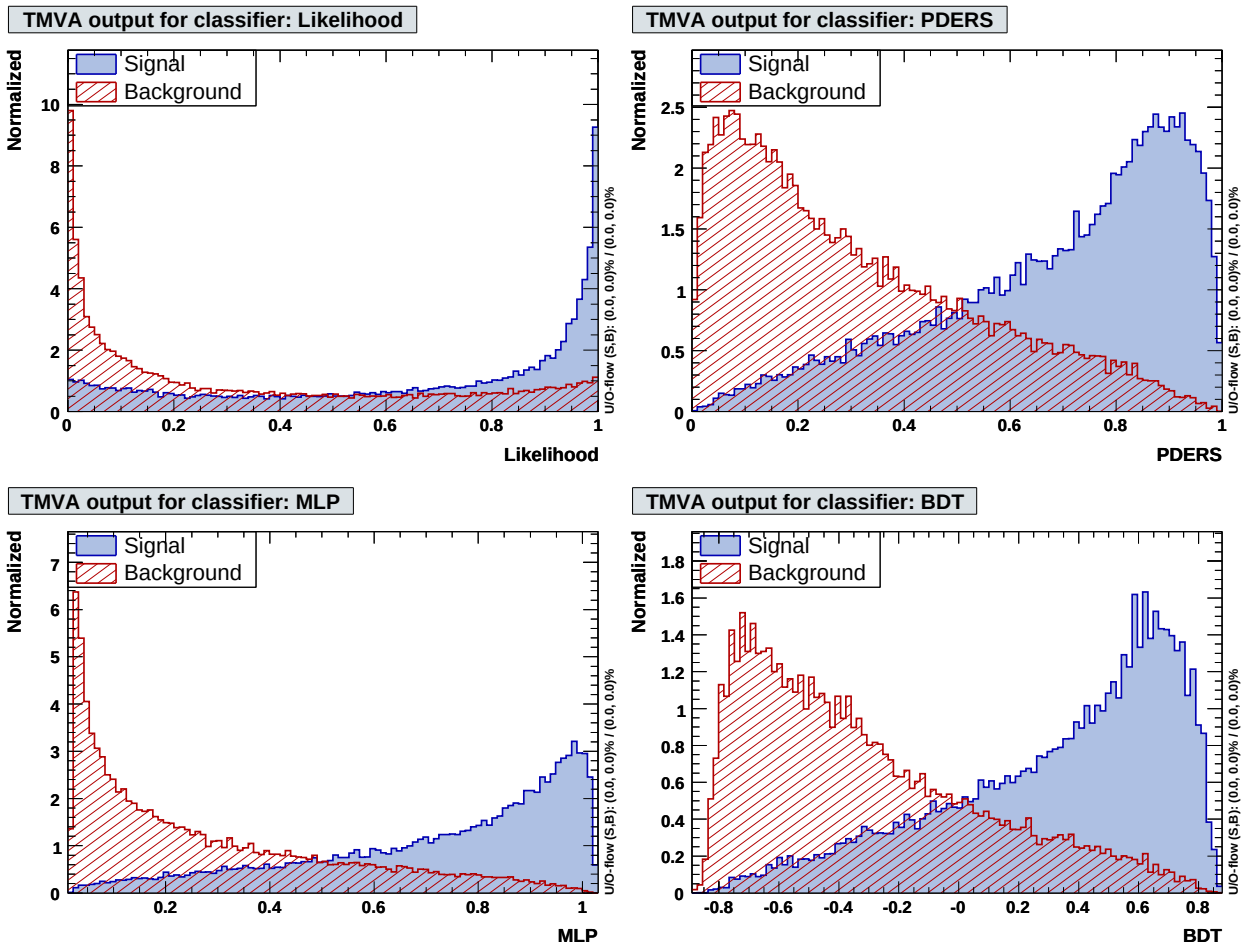

Figure 4: Example plots for classifier output distributions for signal and background events from the academic test sample. Shown are likelihood (upper left), PDE range search (upper right), Multilayer perceptron (MLP – lower left) and boosted decision trees.

message is printed when the option "`H`" is added to the configuration string while booking the method (switch off by setting "`!H`"). The very same help messages are also obtained by clicking the "info" button on the top of the reference tables on the options reference web page: http://tmva.sourceforge.net/optionRef.html.

- The web address of this Users Guide: http://tmva.sourceforge.net/docu/TMVAUsersGuide.pdf.
- The TMVA talk collection: http://tmva.sourceforge.net/talks.shtml.
- TMVA versions in ROOT releases: http://tmva.sourceforge.net/versionRef.html.
- Direct code views via ViewVC: http://tmva.svn.sourceforge.net/viewvc/tmva/trunk/TMVA.
- Class index of TMVA in ROOT: http://root.cern.ch/root/htmldoc/TMVA_Index.html.

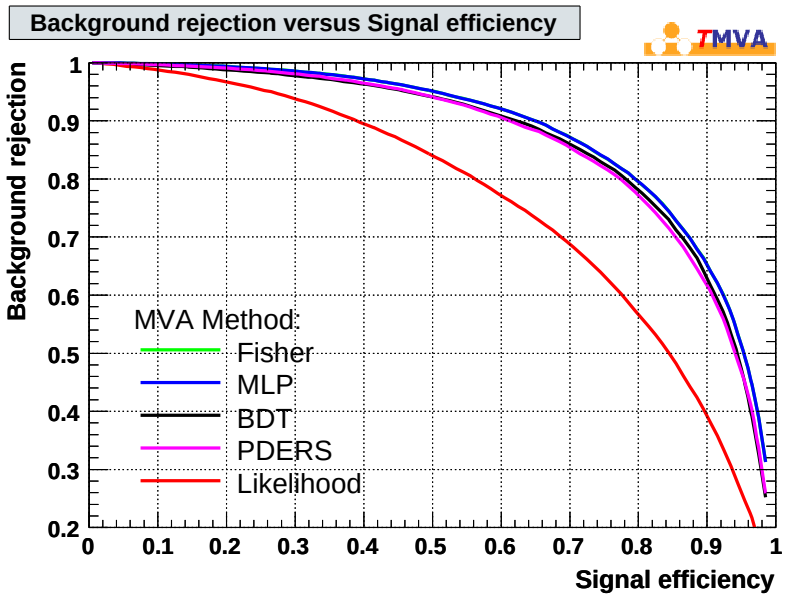

Figure 5: Example for the background rejection versus signal efficiency ("ROC curve") obtained by cutting on the classifier outputs for the events of the test sample.

- Please send questions and/or report problems to the *tmva-users* mailing list: http://sourceforge.net/mailarchive/forum.php?forum_name=tmva-users (posting messages requires prior subscription: https://lists.sourceforge.net/lists/listinfo/tmva-users).

3 Using TMVA

A typical TMVA classification or regression analysis consists of two independent phases: the *training* phase, where the multivariate methods are trained, tested and evaluated, and an *application* phase, where the chosen methods are applied to the concrete classification or regression problem they have been trained for. An overview of the code flow for these two phases as implemented in the examples `TMVAClassification.C` and `TMVAClassificationApplication.C` (for classification – see Sec. 2.5), and `TMVARegression.C` and `TMVARegressionApplication.C` (for regression) are sketched in Fig. 7.

In the training phase, the communication of the user with the data sets and the MVA methods is performed via a `Factory` object, created at the beginning of the program. The TMVA Factory provides member functions to specify the training and test data sets, to register the discriminating input (and – in case of regression – target) variables, and to book the multivariate methods. Subsequently the Factory calls for training, testing and the evaluation of the booked MVA methods. Specific result ("weight") files are created after the training phase by each booked MVA method.

The application of training results to a data set with unknown sample composition (classification) / target value (regression) is governed by the `Reader` object. During initialisation, the user registers

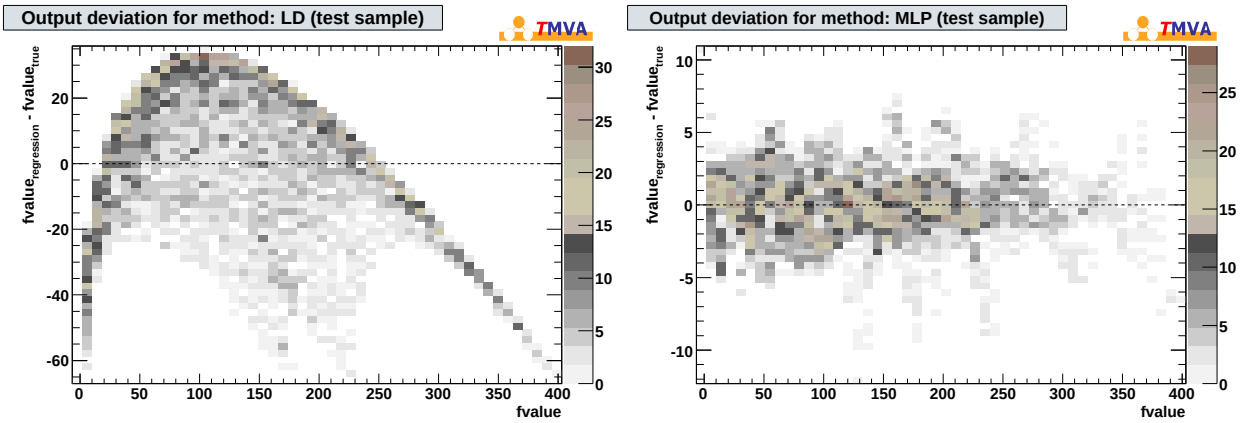

Figure 6: Example plots for the deviation between regression output and target values for a Linear Discriminant (LD – left) and MLP (right). The dependence of the input variables on the target being strongly nonlinear, LD cannot appropriately solve the regression problem.

the input variables[3] together with their local memory addresses, and books the MVA methods that were found to be the most appropriate after evaluating the training results. As booking argument, the name of the weight file is given. The weight file provides for each of the methods full and consistent configuration according to the training setup and results. Within the event loop, the input variables are updated for each event, and the MVA response values and, in some cases, errors are computed.

For standalone use of the trained MVA methods, TMVA also generates lightweight C++ response classes, which contain the encoded information from the weight files so that these are not required anymore (cf. Sec. 3.4).

3.1 The TMVA Factory

The TMVA training phase begins by instantiating a `Factory` object with configuration options listed in Option-Table 1.

```
TMVA::Factory* factory
           = new TMVA::Factory( "<JobName>", outputFile, "<options>" );
```

Code Example 6: Instantiating a Factory class object. The first argument is the user-defined job name that will reappear in the names of the weight files containing the training results. The second argument is the pointer to a writable `TFile` output file created by the user, where control and performance histograms are stored.

[3] This somewhat redundant operation is required to verify the correspondence between the Reader analysis and the weight files used.

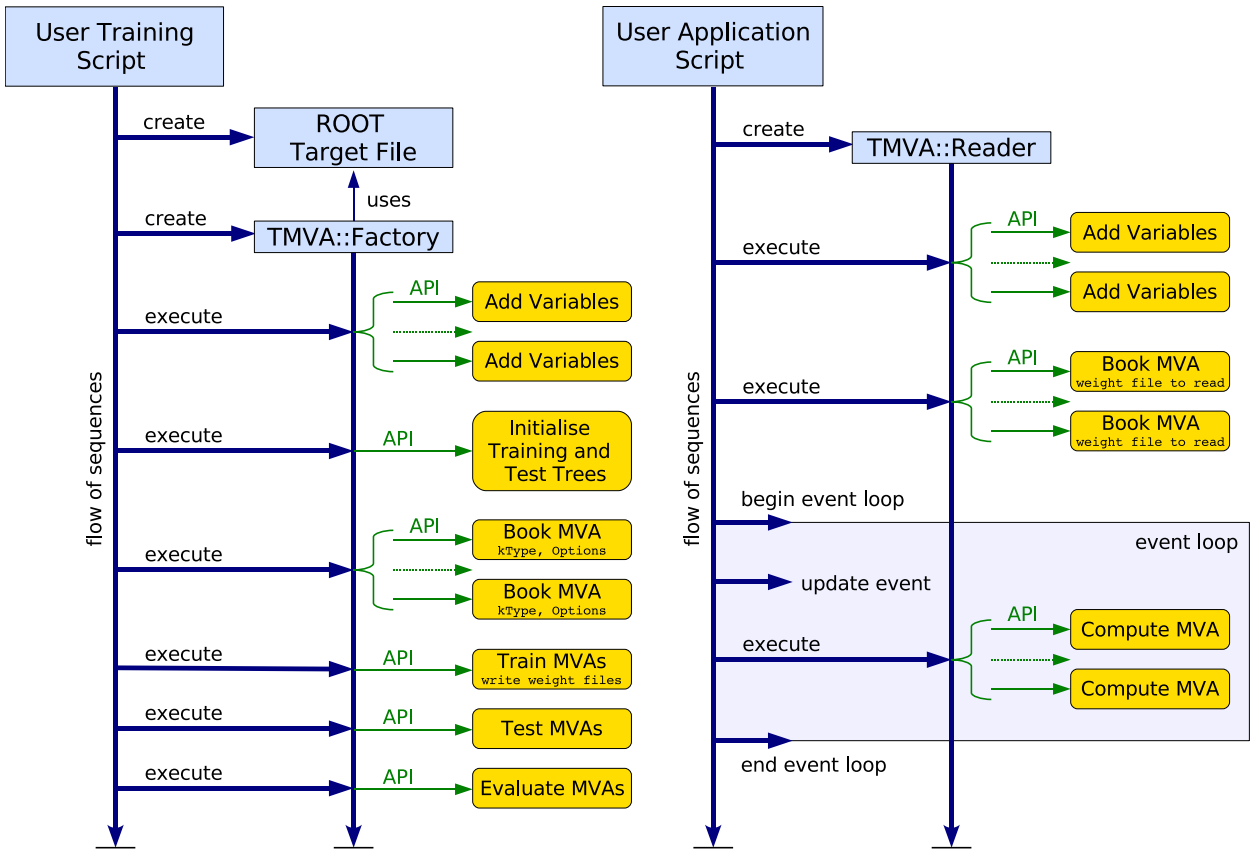

Figure 7: Left: Flow (top to bottom) of a typical TMVA training application. The user script can be a ROOT macro, C++ executable, python script or similar. The user creates a ROOT `TFile`, which is used by the TMVA Factory to store output histograms and trees. After creation by the user, the Factory organises the user's interaction with the TMVA modules. It is the only TMVA object directly created and owned by the user. First the discriminating variables that must be `TFormula`-compliant functions of branches in the training trees are registered. For regression also the target variable must be specified. Then, selected MVA methods are booked through a type identifier and a user-defined unique name, and configuration options are specified via an option string. The TMVA analysis proceeds by consecutively calling the training, testing and performance evaluation methods of the Factory. The training results for all booked methods are written to custom weight files in XML format and the evaluation histograms are stored in the output file. They can be analysed with specific macros that come with TMVA (cf. Tables 2 and 4).
Right: Flow (top to bottom) of a typical TMVA analysis application. The MVA methods qualified by the preceding training and evaluation step are now used to classify data of unknown signal and background composition or to predict a regression target. First, a `Reader` class object is created, which serves as interface to the method's response, just as was the Factory for the training and performance evaluation. The discriminating variables and references to locally declared memory placeholders are registered with the Reader. The variable names and types must be equal to those used for the training. The selected MVA methods are booked with their weight files in the argument, which fully configures them. The user then runs the event loop, where for each event the values of the input variables are copied to the reserved memory addresses, and the MVA response values (and in some cases errors) are computed.

Option	Array	Default	Predefined Values	Description
`V`	–	`False`	–	Verbose flag
`Color`	–	`True`	–	Flag for coloured screen output (default: True, if in batch mode: False)
`Transformations`	–		–	List of transformations to test; formatting example: Transformations=I;D;P;G,D, for identity, decorrelation, PCA, and Gaussianisation followed by decorrelation transformations
`Silent`	–	`False`	–	Batch mode: boolean silent flag inhibiting any output from TMVA after the creation of the factory class object (default: False)
`DrawProgressBar`	–	`True`	–	Draw progress bar to display training, testing and evaluation schedule (default: True)

Option Table 1: Configuration options reference for class: *Factory*. Coloured output is switched on by default, except when running ROOT in batch mode (i.e., when the '`-b`' option of the CINT interpreter is invoked). The list of transformations contains a default set of data preprocessing steps for test and visualisation purposes only. The usage of preprocessing transformations in conjunction with MVA methods must be configured when booking the methods.

3.1.1 Specifying training and test data

The input data sets used for training and testing of the multivariate methods need to be handed to the Factory. TMVA supports ROOT `TTree` and derived `TChain` objects as well as text files. If ROOT trees are used for classification problems, the signal and background events can be located in the same or in different trees. Overall weights can be specified for the signal and background training data (the treatment of event-by-event weights is discussed below).

Specifying **classification training data** in ROOT tree format with signal and background events being located in different trees:

```
// Get the signal and background trees from TFile source(s);
// multiple trees can be registered with the Factory
TTree* sigTree  = (TTree*)sigSrc->Get( "<YourSignalTreeName>"   );
TTree* bkgTreeA = (TTree*)bkgSrc->Get( "<YourBackgrTreeName_A>" );
TTree* bkgTreeB = (TTree*)bkgSrc->Get( "<YourBackgrTreeName_B>" );
TTree* bkgTreeC = (TTree*)bkgSrc->Get( "<YourBackgrTreeName_C>" );

// Set the event weights per tree (these weights are applied in
// addition to individual event weights that can be specified)
Double_t sigWeight  = 1.0;
Double_t bkgWeightA = 1.0, bkgWeightB = 0.5, bkgWeightC = 2.0;

// Register the trees
factory->AddSignalTree    ( sigTree,  sigWeight  );
factory->AddBackgroundTree( bkgTreeA, bkgWeightA );
factory->AddBackgroundTree( bkgTreeB, bkgWeightB );
factory->AddBackgroundTree( bkgTreeC, bkgWeightC );
```

Code Example 7: Registration of signal and background ROOT trees read from `TFile` sources. Overall signal and background weights per tree can also be specified. The `TTree` object may be replaced by a `TChain`.

Specifying **classification training data** in ROOT tree format with signal and background events being located in the same tree:

```
TTree* inputTree = (TTree*)source->Get( "<YourTreeName>" );

TCut signalCut = ...;  // how to identify signal events
TCut backgrCut = ...;  // how to identify background events

factory->SetInputTrees( inputTree, signalCut, backgrCut );
```

Code Example 8: Registration of a single ROOT tree containing the input data for signal *and* background, read from a `TFile` source. The `TTree` object may be replaced by a `TChain`. The cuts identify the event species.

Specifying **classification training data** in text format:

```
// Text file format (available types: 'F' and 'I')
//   var1/F:var2/F:var3/F:var4/F
//   0.21293  -0.49200  -0.58425  -0.70591
//   ...
TString sigFile = "signal.txt";     // text file for signal
TString bkgFile = "background.txt"; // text file for background

Double_t sigWeight = 1.0; // overall weight for all signal events
Double_t bkgWeight = 1.0; // overall weight for all background events

factory->SetInputTrees( sigFile, bkgFile, sigWeight, bkgWeight );
```

Code Example 9: Registration of signal and background text files. Names and types of the input variables are given in the first line, followed by the values.

Specifying **regression training data** in ROOT tree format:

```
factory->AddRegressionTree( regTree, weight );
```

Code Example 10: Registration of a ROOT tree containing the input and target variables. An overall weight per tree can also be specified. The `TTree` object may be replaced by a `TChain`.

3.1.2 Defining input variables, targets and event weights

The variables in the input trees used to train the MVA methods are registered with the Factory using the `AddVariable` method. It takes the variable name (string), which must have a correspondence in the input ROOT tree or input text file, and optionally a number type (`'F'` (default) and `'I'`). The type is used to inform the method whether a variable takes continuous floating point or discrete values.[4] Note that `'F'` indicates *any* floating point type, i.e., `float` *and* `double`. Correspondingly, `'I'` stands for integer, *including* `int`, `short`, `char`, and the corresponding `unsigned` types. Hence, if a variable in the input tree is `double`, it should be declared `'F'` in the `AddVariable` call.

It is possible to specify variable expressions, just as for the `TTree::Draw` command (the expression is interpreted as a `TTreeFormula`, including the use of arrays). Expressions may be abbreviated for more concise screen output (and plotting) purposes by defining shorthand-notation *labels* via the assignment operator `:=`.

In addition, two more arguments may be inserted into the `AddVariable` call, allowing the user to specify *titles* and *units* for the input variables for displaying purposes.

[4] For example for the projective likelihood method, a histogram out of discrete values would not (and should not) be interpolated between bins.

The following code example revises all possible options to declare an input variable:

```
factory->AddVariable( "<YourDescreteVar>",                  'I' );
factory->AddVariable( "log(<YourFloatingVar>)",             'F' );
factory->AddVariable( "SumLabel := <YourVar1>+<YourVar2>",  'F' );
factory->AddVariable( "<YourVar3>", "Pretty Title", "Unit", 'F' );
```

Code Example 11: Declaration of variables used to train the MVA methods. Each variable is specified by its name in the training tree (or text file), and optionally a type (`'F'` for floating point and `'I'` for integer, `'F'` is default if nothing is given). Note that `'F'` indicates *any* floating point type, i.e., `float` *and* `double`. Correspondingly, `'I'` stands for integer, *including* `int`, `short`, `char`, and the corresponding `unsigned` types. Hence, even if a variable in the input tree is `double`, it should be declared `'F'` here. Here, `YourVar1` has discrete values and is thus declared as an integer. Just as in the `TTree::Draw` command, it is also possible to specify expressions of variables. The `:=` operator defines labels (third row), used for shorthand notation in screen outputs and plots. It is also possible to define titles and units for the variables (fourth row), which are used for plotting. If labels *and* titles are defined, labels are used for abbreviated screen outputs, and titles for plotting.

It is possible to define *spectator variables*, which are part of the input data set, but which are not used in the MVA training, test nor during the evaluation. They are copied into the `TestTree`, together with the used input variables and the MVA response values for each event, where the spectator variables can be used for correlation tests or others. Spectator variables are declared as follows:

```
factory->AddSpectator( "<YourSpectatorVariable>" );
factory->AddSpectator( "log(<YourSpectatorVariable>)" );
factory->AddSpectator( "<YourSpectatorVariable>", "Pretty Title", "Unit" );
```

Code Example 12: Various ways to declare a spectator variable, not participating in the MVA anlaysis, but written into the final `TestTree`.

For a regression problem, the target variable is defined similarly, without however specifying a number type:

```
factory->AddTarget( "<YourRegressionTarget1>" );
factory->AddTarget( "log(<YourRegressionTarget2>)" );
factory->AddTarget( "<YourRegressionTarget3>", "Pretty Title", "Unit" );
```

Code Example 13: Various ways to declare the target variables used to train a multivariate regression method. If the MVA method supports multi-target (multidimensional) regression, more than one regression target can be defined.

Individual events can be weighted, with the weights being a column or a function of columns of the

input data sets. To specify the weights to be used for the training use the command:

```
factory->SetWeightExpression( "<YourWeightExpression>" );
```

Code Example 14: Specification of individual weights for the training events. The expression must be a function of variables present in the input data set.

3.1.3 Negative event weights

In next-to-leading order Monte Carlo generators, events with (unphysical) negative weights may occur in some phase space regions. Such events are often troublesome to deal with, and it depends on the concrete implementation of the MVA method, whether or not they are treated properly. Among those methods that correctly incorporate events with negative weights are likelihood and multi-dimensional probability density estimators, but also decision trees. A summary of this feature for all TMVA methods is given in Table 7. In cases where a method does *not* properly treat events with negative weights, it is advisable to ignore such events for the training - but to include them in the performance evaluation to not bias the results. This can be explicitly requested for each MVA method via the boolean configuration option `IgnoreNegWeightsInTraining` (cf. Option Table 9 on page 58).

3.1.4 Preparing the training and test data

The input events that are handed to the Factory are internally copied and split into one *training* and one *test* ROOT tree. This guarantees a statistically independent evaluation of the MVA algorithms based on the test sample.[5] The numbers of events used in both samples are specified by the user. They must not exceed the entries of the input data sets. In case the user has provided a ROOT tree, the event copy can (and should) be accelerated by disabling all branches not used by the input variables.

It is possible to apply selection requirements (cuts) upon the input events. These requirements can depend on any variable present in the input data sets, i.e., they are not restricted to the variables used by the methods. The full command is as follows:

[5] A fully unbiased training and evaluation requires at least three statistically independent data sets. See comments in Footnote 9 on page 27.

```
TCut preselectionCut = "<YourSelectionString>";
factory->PrepareTrainingAndTestTree( preselectionCut, "<options>" );
```

Code Example 15: Preparation of the internal TMVA training and test trees. The sizes (number of events) of these trees are specified in the configuration option string. For classification problems, they can be set individually for signal and background. Note that the preselection cuts are applied before the training and test samples are created, i.e., the tree sizes apply to numbers of *selected* events. It is also possible to choose among different methods to select the events entering the training and test trees from the source trees. All options are described in Option-Table 2. See also the text for further information.

For **classification**, the numbers of signal and background events used for training and testing are specified in the configuration string by the variables `nTrain_Signal`, `nTrain_Background`, `nTest_Signal` and `nTest_Background` (for example, `"nTrain_Signal=5000:nTrain_Background=5000:nTest_Signal=4000:nTest_Background=5000"`). The default value (zero) signifies that all available events are taken, e.g., if `nTrain_Signal=5000` and `nTest_Signal=0`, and if the total signal sample has 15000 events, then 5000 signal events are used for training and the remaining 10000 events are used for testing. If `nTrain_Signal=0` and `nTest_Signal=0`, the signal sample is split in half for training and testing. The same rules apply to background. Since zero is default, not specifying anything corresponds to splitting the samples in two halves.

For **regression**, only the sizes of the train and test samples are given, e.g., `"nTrain_Regression=0:nTest_Regression=0"`, so that one half of the input sample is used for training and the other half for testing.

The option `SplitMode` defines how the training and test samples are selected from the source trees. With `SplitMode=Random`, events are selected randomly. With `SplitMode=Alternate`, events are chosen in alternating turns for the training and test samples as they occur in the source trees until the desired numbers of training and test events are selected. In the `SplitMode=Block` mode the first `nTrain_Signal` and `nTrain_Background` (classification), or `nTrain_Regression` events (regression) of the input data set are selected for the training sample, and the next `nTest_Signal` and `nTest_Background` or `nTest_Regression` events comprise the test data. This is usually not desired for data that contains varying conditions over the range of the data set. For the `Random` selection mode, the seed of the random generator can be set. With `SplitSeed=0` the generator returns a different random number series every time. The default seed of 100 results in the same training and test samples each time TMVA is run (as does any other seed apart from 0).

In some cases event weights are given by Monte Carlo generators, and may turn out to be overall very small or large numbers. To avoid artifacts due to this, TMVA internally renormalises the signal and background weights so that their sums over all events equal the respective numbers of events in the two samples. The renormalisation is optional and can be modified with the configuration option `NormMode` (cf. Table 2). Possible settings are: `None`: no renormalisation is applied (the weights are used as given), `NumEvents` (default): renormalisation to sums of events as described above, `EqualNumEvents`: the event weights are renormalised so that both, the sum of signal and the sum of background weights equal the number of signal events in the sample.

Option	Array	Default	Predefined Values	Description
SplitMode	–	Random	Random, Alternate, Block	Method of picking training and testing events (default: random)
SplitSeed	–	100	–	Seed for random event shuffling
NormMode	–	NumEvents	None, NumEvents, EqualNumEvents	Overall renormalisation of event-by-event weights (NumEvents: average weight of 1 per event, independently for signal and background; EqualNumEvents: average weight of 1 per event for signal, and sum of weights for background equal to sum of weights for signal)
nTrain_Signal	–	0	–	Number of training events of class Signal (default: 0 = all)
nTest_Signal	–	0	–	Number of test events of class Signal (default: 0 = all)
nTrain_Background	–	0	–	Number of training events of class Background (default: 0 = all)
nTest_Background	–	0	–	Number of test events of class Background (default: 0 = all)
V	–	False	–	Verbosity (default: true)
VerboseLevel	–	Info	Debug, Verbose, Info	VerboseLevel (Debug/Verbose/Info)

Option Table 2: Configuration options reference in call `Factory::PrepareTrainingAndTestTree(..)`. For regression, `nTrain_Signal` and `nTest_Signal` are replaced by `nTrain_Regression` and `nTest_Regression`, respectively, and `nTrain_Background` and `nTest_Background` are removed. See also Code-Example 15 and comments in the text.

3.1.5 Booking MVA methods

All MVA methods are booked via the Factory by specifying the method's type, plus a unique name chosen by the user, and a set of specific configuration options encoded in a string qualifier.[6] If the same method type is booked several times with different options (which is useful to compare different sets of configurations for optimisation purposes), the specified names must be different to distinguish the instances and their weight files. A booking example for the likelihood method is given in Code Example 16 below. Detailed descriptions of the configuration options are given in the corresponding tools and MVA sections of this Users Guide, and booking examples for most of the methods are given in Appendix A. With the MVA booking the initialisation of the Factory is complete and no MVA-specific actions are left to do. The Factory takes care of the subsequent training, testing and evaluation of the MVA methods.

[6] In the TMVA package all MVA methods are derived from the abstract interface `IMethod` and the base class `MethodBase`.

```
factory->BookMethod( TMVA::Types::kLikelihood, "LikelihoodD",
                     "!H:!V:!TransformOutput:PDFInterpol=Spline2:\
                      NSmoothSig[0]=20:NSmoothBkg[0]=20:NSmooth=5:\
                      NAvEvtPerBin=50:VarTransform=Decorrelate" );
```

Code Example 16: Example booking of the likelihood method. The first argument is a unique type enumerator (the available types can be looked up in `src/Types.h`), the second is a user-defined name which must be unique among all booked MVA methods, and the third is a configuration option string that is specific to the method. For options that are not explicitly set in the string default values are used, which are printed to standard output. The syntax of the options should be explicit from the above example. Individual options are separated by a ':'. Boolean variables can be set either explicitly as `MyBoolVar=True/False`, or just via `MyBoolVar/!MyBoolVar`. All specific options are explained in the tools and MVA sections of this Users Guide. There is no difference in the booking of methods for classification or regression applications. See Appendix A on page 134 for a complete booking list of all MVA methods in TMVA.

3.1.6 Help option for MVA booking

Upon request via the configuration option "`H`" (see code example above) the TMVA methods print concise help messages. These include a brief description of the algorithm, a performance assessment, and hints for setting the most important configuration options. The messages can also be evoked by the command `factory->PrintHelpMessage("<MethodName>")`.

3.1.7 Training the MVA methods

The training of the booked methods is invoked by the command:

```
factory->TrainAllMethods();
```

Code Example 17: Executing the MVA training via the Factory.

The training results are stored in the weight files which are saved in the directory `weights` (which, if not existing is created).[7] The weight files are named `Jobname_MethodName.weights.`⟨`extension`⟩, where the job name has been specified at the instantiation of the Factory, and `MethodName` is the unique method name specified in the booking command. Each method writes a custom weight file in XML format (extension is `xml`), where the configuration options, controls and training results for the method are stored.

[7] The default weight file directory name can be modified from the user script through the global configuration variable `(TMVA::gConfig().GetIONames()).fWeightFileDir`.

3.1.8 Testing the MVA methods

The trained MVA methods are applied to the test data set and provide scalar outputs according to which an event can be classified as either signal or background, or which estimate the regression target.[8] The MVA outputs are stored in the test tree (`TestTree`) to which a column is added for each booked method. The tree is eventually written to the output file and can be directly analysed in a ROOT session. The testing of all booked methods is invoked by the command:

```
factory->TestAllMethods();
```

Code Example 18: Executing the validation (testing) of the MVA methods via the Factory.

3.1.9 Evaluating the MVA methods

The Factory and data set classes of TMVA perform a preliminary property assessment of the input variables used by the MVA methods, such as computing correlation coefficients and ranking the variables according to their separation (for classification), or according to their correlations with the target variable(s) (for regression). The results are printed to standard output.

The performance evaluation in terms of signal efficiency, background rejection, faithful estimation of a regression target, etc., of the trained and tested MVA methods is invoked by the command:

```
factory->EvaluateAllMethods();
```

Code Example 19: Executing the performance evaluation via the Factory.

The performance measures differ between classification and regression problems. They are summarised below.

3.1.10 Classification performance evaluation

After training and testing, the linear correlation coefficients among the classifier outputs are printed. In addition, overlap matrices are derived (and printed) for signal and background that determine the fractions of signal and background events that are equally classified by each pair of classifiers. This is useful when two classifiers have similar performance, but a significant fraction of non-overlapping

[8] In classification mode, TMVA discriminates signal from background in data sets with unknown composition of these two samples. In frequent use cases the background (sometimes also the signal) consists of a variety of different populations with characteristic properties, which could call for classifiers with more than two discrimination classes. However, in practise it is usually possible to serialise background fighting by training individual classifiers for each background source, and applying consecutive requirements to these. Since TMVA 4, the framework supports multi-class classification. However, the individual MVA methods have not yet been prepared for it.

events. In such a case a combination of the classifiers (e.g., in a *Committee* classifier) could improve the performance (this can be extended to any combination of any number of classifiers).

The optimal method to be used for a specific analysis strongly depends on the problem at hand and no general recommendations can be given. To ease the choice TMVA computes a number of benchmark quantities that assess the performance of the methods on the independent test sample. For classification these are

- The **signal efficiency at three representative background efficiencies** (the efficiency is equal to 1 − rejection) obtained from a cut on the classifier output. Also given is the area of the background rejection versus signal efficiency function (the larger the area the better the performance).
- The **separation** $\langle S^2 \rangle$ of a classifier y, defined by the integral [7]

$$\langle S^2 \rangle = \frac{1}{2} \int \frac{(\hat{y}_S(y) - \hat{y}_B(y))^2}{\hat{y}_S(y) + \hat{y}_B(y)} dy \,, \tag{1}$$

  where $\hat{y}_S$ and $\hat{y}_B$ are the signal and background PDFs of y, respectively (cf. Sec. 3.1.13). The separation is zero for identical signal and background shapes, and it is one for shapes with no overlap.
- The discrimination **significance** of a classifier, defined by the difference between the classifier means for signal and background divided by the quadratic sum of their root-mean-squares.

The results of the evaluation are printed to standard output. Smooth background rejection/efficiency versus signal efficiency curves are written to the output ROOT file, and can be plotted using custom macros (see Sec. 3.2).

3.1.11 Regression performance evaluation

Ranking for regression is based on the correlation strength between the input variables or MVA method response and the regression target. Several correlation measures are implemented in TMVA to capture and quantify nonlinear dependencies. Their results are printed to standard output.

- The **Correlation** between two random variables X and Y is usually measured with the correlation coefficient ρ, defined by

$$\rho(X,Y) = \frac{\mathrm{cov}(X,Y)}{\sigma_X \sigma_Y}. \tag{2}$$

  The correlation coefficient is symmetric in X and Y, lies within the interval $[-1,1]$, and quantifies by definition a linear relationship. Thus $\rho = 0$ holds for independent variables, but the converse is not true in general. In particular, higher order functional or non-functional relationships may not, or only marginally, be reflected in the value of ρ (see Fig. 8).

- The **correlation ratio** is defined by

$$\eta^2(Y|X) = \frac{\sigma_{E(Y|X)}}{\sigma_Y}\,, \tag{3}$$

where

$$E(Y|X) = \int y\ P(y|x)\ dy\,, \tag{4}$$

is the conditional expectation of Y given X with the associated conditional probability density function $P(Y|X)$. The correlation ratio η^2 is in general not symmetric and its value lies within $[0,1]$, according to how well the data points can be fitted with a linear or nonlinear regression curve. Thus non-functional correlations cannot be accounted for by the correlation ratio. The following relations can be derived for η^2 and the squared correlation coefficient ρ^2 [9]:

  - $\rho^2 = \eta^2 = 1$, if X and Y are in a strict linear functional relationship.
  - $\rho^2 \leq \eta^2 = 1$, if X and Y are in a strict nonlinear functional relationship.
  - $\rho^2 = \eta^2 < 1$, if there is no strict functional relationship but the regression of X on Y is exactly linear.
  - $\rho^2 < \eta^2 < 1$, if there is no strict functional relationship but some nonlinear regression curve is a better fit then the best linear fit.

Some characteristic examples and their corresponding values for η^2 are shown in Fig. 8. In the special case, where all data points take the same value, η is undefined.

- **Mutual information** allows to detect any predictable relationship between two random variables, be it of functional or non-functional form. It is defined by [10]

$$I(X,Y) = \sum_{X,Y} P(X,Y) \ln \frac{P(X,Y)}{P(X)P(Y)}\,, \tag{5}$$

where $P(X,Y)$ is the joint probability density function of the random variables X and Y, and $P(X)$, $P(Y)$ are the corresponding marginal probabilities. Mutual information originates from information theory and is closely related to entropy which is a measure of the uncertainty associated with a random variable. It is defined by

$$H(X) = -\sum_{X} P(X) \ln P(X)\,, \tag{6}$$

where X is the discrete random variable and $P(X)$ the associated probability density function.

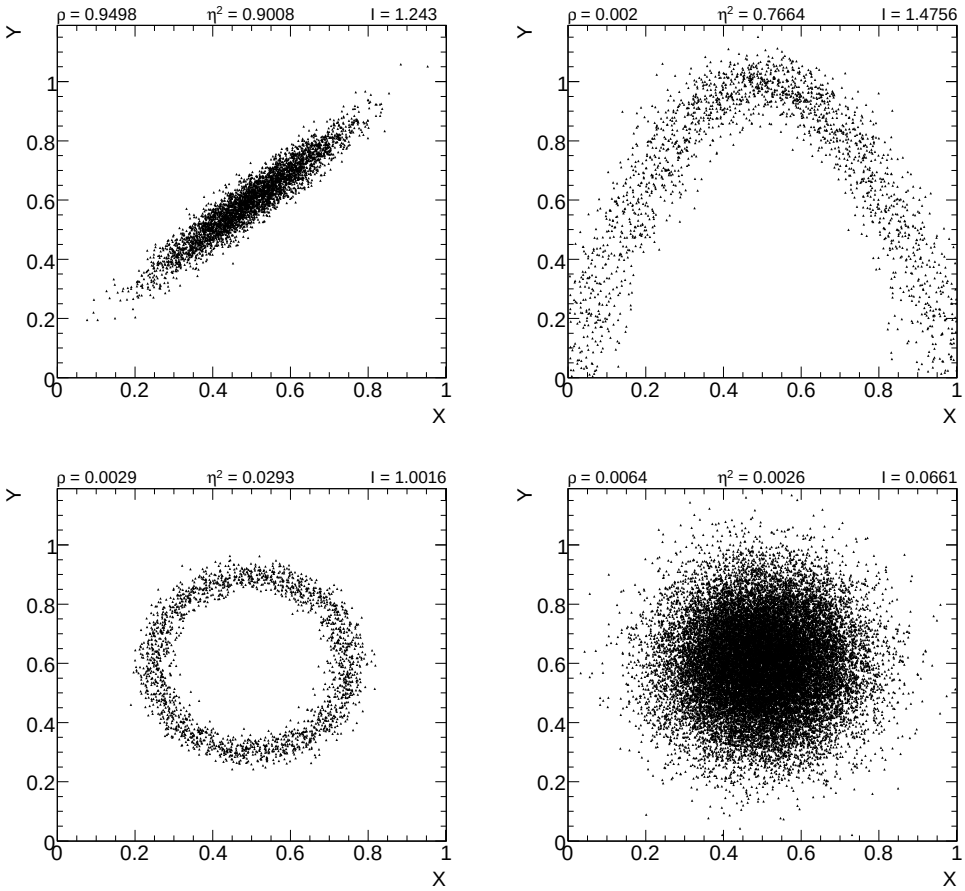

Figure 8: Various types of correlations between two random variables and their corresponding values for the correlation coefficient ρ, the correlation ratio η, and mutual information I. Linear relationship (upper left), functional relationship (upper right), non-functional relationship (lower left), and independent variables (lower right).

The connection between the two quantities is given by the following transformation

$$I(X,Y) = \sum_{X,Y} P(X,Y) \ln \frac{P(X,Y)}{P(X)P(Y)} \tag{7}$$

$$= \sum_{X,Y} P(X,Y) \ln \frac{P(X|Y)}{P_X(X)} \tag{8}$$

$$= -\sum_{X,Y} P(X,Y) \ln P(X) + \sum_{X,Y} P(X,Y) \ln P(X|Y) \tag{9}$$

$$= -\sum_{X,Y} P(X) \ln P(X) - (-\sum_{X,Y} P(X,Y) \ln P(X|Y)) \tag{10}$$

$$= H(X) - H(X|Y)\,, \tag{11}$$

where $H(X|Y)$ is the conditional entropy of X given Y. Thus mutual information is the reduction of the uncertainty in variable X due to the knowledge of Y. Mutual information

ρ_{PDF}	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.9999
ρ	0.006	0.092	0.191	0.291	0.391	0.492	0.592	0.694	0.795	0.898	1.0
η^2	0.004	0.012	0.041	0.089	0.156	0.245	0.354	0.484	0.634	0.806	1.0
I	0.093	0.099	0.112	0.139	0.171	0.222	0.295	0.398	0.56	0.861	3.071

Table 1: Comparison of the correlation coefficient ρ, correlation ratio η, and mutual information I for two-dimensional Gaussian toy Monte-Carlo distributions with linear correlations as indicated (20000 data points/100×100 bins .

is symmetric and takes positive absolute values. In the case of two completely independent variables $I(X, Y)$ is zero.

For experimental measurements the joint and marginal probability density functions are a priori unknown and must be approximated by choosing suitable binning procedures such as kernel estimation techniques (see, e.g., [11]). Consequently, the values of $I(X, Y)$ for a given data set will strongly depend on the statistical power of the sample and the chosen binning parameters.

For the purpose of ranking variables from data sets of equal statistical power and identical binning, however, we assume that the evaluation from a simple two-dimensional histogram without further smoothing is sufficient.

A comparison of the correlation coefficient ρ, the correlation ratio η, and mutual information I for linearly correlated two-dimensional Gaussian toy MC simulations is shown in Table 1.

3.1.12 Overtraining

Overtraining occurs when a machine learning problem has too few degrees of freedom, because too many model parameters of an algorithm were adjusted to too few data points. The sensitivity to overtraining therefore depends on the MVA method. For example, a Fisher (or *linear*) discriminant can hardly ever be overtrained, whereas, without the appropriate counter measures, boosted decision trees usually suffer from at least partial overtraining, owing to their large number of nodes. Overtraining leads to a seeming increase in the classification or regression performance over the objectively achievable one, if measured on the training sample, and to an effective performance decrease when measured with an independent test sample. A convenient way to detect overtraining and to measure its impact is therefore to compare the performance results between training and test samples. Such a test is performed by TMVA with the results printed to standard output.

Various method-specific solutions to counteract overtraining exist. For example, binned likelihood reference distributions are smoothed before interpolating their shapes, or unbinned kernel density estimators smear each training event before computing the PDF; neural networks steadily monitor the convergence of the error estimator between training and test samples[9] suspending the training

[9] Proper training and validation requires three statistically independent data sets: one for the parameter optimi-

when the test sample has passed its minimum; the number of nodes in boosted decision trees can be reduced by removing insignificant ones ("tree pruning"), etc.

3.1.13 Other representations of MVA outputs for classification: probabilities and *Rarity*

In addition to the MVA response value y of a classifier, which is typically used to place a cut for the classification of an event as either signal or background, or which could be used in a subsequent likelihood fit, TMVA also provides the classifier's signal and background PDFs, $\hat{y}_{S(B)}$. The PDFs can be used to derive classification probabilities for individual events, or to compute any kind of transformation of which the *Rarity* transformation is implemented in TMVA.

- **Classification probability**: The techniques used to estimate the shapes of the PDFs are those developed for the likelihood classifier (see Sec. 8.2.2 for details) and can be customised individually for each method (the control options are given in Sec. 8). The probability for event i to be of signal type is given by,

$$P_S(i) = \frac{f_S \cdot \hat{y}_S(i)}{f_S \cdot \hat{y}_S(i) + (1 - f_S) \cdot \hat{y}_B(i)} \,, \tag{12}$$

  where $f_S = N_S/(N_S + N_B)$ is the expected signal fraction, and $N_{S(B)}$ is the expected number of signal (background) events (default is $f_S = 0.5$).[10]

- **Rarity**: The Rarity $\mathcal{R}(y)$ of a classifier y is given by the integral [8]

$$\mathcal{R}(y) = \int_{-\infty}^{y} \hat{y}_B(y')\, dy' \,, \tag{13}$$

  which is defined such that $\mathcal{R}(y_B)$ for background events is uniformly distributed between 0 and 1, while signal events cluster towards 1. The signal distributions can thus be directly compared among the various classifiers. The stronger the peak towards 1, the better is the discrimination. Another useful aspect of the Rarity is the possibility to directly visualise deviations of a test background (which could be physics data) from the training sample, by exhibition of non-uniformity.

  The Rarity distributions of the Likelihood and Fisher classifiers for the example used in Sec. 2 are plotted in Fig. 9. Since Fisher performs better (cf. Fig. 5 on page 12), its signal distribution is stronger peaked towards 1. By construction, the background distributions are uniform within statistical fluctuations.

The probability and Rarity distributions can be plotted with dedicated macros, invoked through corresponding GUI buttons.

sation, another one for the overtraining detection, and the last one for the performance validation. In TMVA, the last two samples have been merged to increase statistics. The (usually insignificant) bias introduced by this on the evaluation results does not affect the analysis as far as classification cut efficiencies or the regression resolution are independently validated with data.

[10] The P_S distributions may exhibit a somewhat peculiar structure with frequent narrow peaks. They are generated by regions of classifier output values in which $\hat{y}_S \propto \hat{y}_B$ for which P_S becomes a constant.

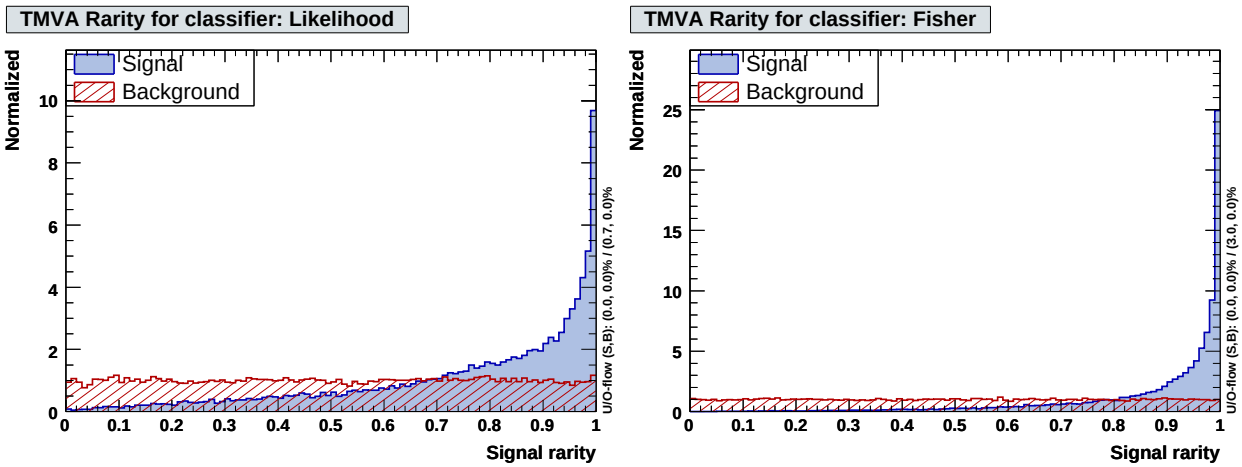

Figure 9: Example plots for classifier Rarity distributions for signal and background events from the academic test sample. Shown are likelihood (left) and Fisher (right).

3.2 ROOT macros to plot training, testing and evaluation results

TMVA provides simple GUIs (`TMVAGui.C` and `TMVARegGui.C`, see Fig. 1), which interface ROOT macros that visualise the various steps of the training analysis. The macros are respectively located in `TMVA/macros/` (Sourceforge.net distribution) and `$ROOTSYS/tmva/test/` (ROOT distribution), and can also be executed from the command line. They are described in Tables 2 and 4. All plots drawn are saved as *png* files (or optionally as *eps*, *gif* files) in the macro subdirectory `plots` which, if not existing, is created.

The binning and histogram boundaries for some of the histograms created during the training, testing and evaluation phases are controlled via the global singleton class `TMVA::Config`. They can be modified as follows:

```
// Modify settings for the variable plotting
(TMVA::gConfig().GetVariablePlotting()).fTimesRMS = 8.0;
(TMVA::gConfig().GetVariablePlotting()).fNbins1D  = 60.0;
(TMVA::gConfig().GetVariablePlotting()).fNbins2D  = 300.0;

// Modify the binning in the ROC curve (for classification only)
(TMVA::gConfig().GetVariablePlotting()).fNbinsXOfROCCurve = 100;

// For file name settings, modify the struct TMVA::Config::IONames
(TMVA::gConfig().GetIONames()).fWeightFileDir = "myWeightFileDir";
```

Code Example 20: Modifying global parameter settings for the plotting of the discriminating input variables. The values given are the TMVA defaults. Consult the class files Config.h and Config.cxx for all available global configuration variables and their default settings, respectively. Note that the additional parentheses are mandatory when used in CINT.

Macro	Description
`variables.C`	Plots the signal and background MVA input variables (training sample). The second argument sets the directory, which determines the preprocessing type (`InputVariables_Id` for default identity transformation, cf. Sec. 4.1). The third argument is a title, and the fourth argument is a flag whether or not the input variables served a regression analysis.
`correlationscatter.C`	Plots superimposed scatters and profiles for all pairs of input variables used during the training phase (separate plots for signal and background in case of classification). The arguments are as above.
`correlations.C`	Plots the linear correlation matrices for the input variables in the training sample (distinguishing signal and background for classification).
`mvas.C`	Plots the classifier response distributions of the test sample for signal and background. The second argument (`HistType=0,1,2,3`) allows to also plot the probability (1) and Rarity (2) distributions of the classifiers, as well as a comparison of the output distributions between test and training samples. Plotting of probability and Rarity requires the `CreateMVAPdfs` option for the classifier to be set to true.
`mvaeffs.C`	Signal and background efficiencies, obtained from cutting on the classifier outputs, versus the cut value. Also shown are the signal purity and the signal efficiency times signal purity corresponding to the expected number of signal and background events before cutting (numbers given by user). The optimal cuts according to the best significance are printed on standard output.
`efficiencies.C`	Background rejection (second argument `type=2`, default), or background efficiency (`type=1`), versus signal efficiency for the classifiers (test sample). The efficiencies are obtained by cutting on the classifier outputs. This is traditionally the best plot to assess the overall discrimination performance (ROC curve).
`paracoor.C`	Draws diagrams of "Parallel coordinates" [31] for signal and background, used to visualise the correlations among the input variables, but also between the MVA output and input variables (indicating the importance of the variables).

Table 2: ROOT macros for the representation of the TMVA input variables and **classification results**. All macros take as first argument the name of the ROOT file containing the histograms (default is `TMVA.root`). They are conveniently called via the `TMVAGui.C` GUI (the first three macros are also called from the regression GUI `TMVARegGui.C`). Macros for the representation of regression results are given in Table 3. Plotting macros for MVA method specific information are listed in Table 4.

Macro	Description
`deviations.C`	Plots the linear deviation between regression target value and MVA response or input variables for test and training samples.
`regression_averagedevs.C`	Draws the average deviation between the MVA output and the regression target value for all trained methods.

Table 3: ROOT macros for the representation of the TMVA **regression results**. All macros take as first argument the name of the ROOT file containing the histograms (default is `TMVA.root`). They are conveniently called from the `TMVARegGui.C` GUI.

3.3 The TMVA Reader

After training and evaluation, the most performing MVA methods are chosen and used to classify events in data samples with unknown signal and background composition, or to predict values of a regression target. An example of how this *application phase* is carried out is given in `TMVA/macros/TMVAClassificationApplication.C` and `TMVA/macros/TMVARegressionApplication.C` (Sourceforge.net), or `$ROOTSYS/tmva/test/TMVAClassificationApplication.C` and `$ROOTSYS/tmva/test/TMVARegressionApplication.C` (ROOT).

Analogously to the Factory, the communication between the user application and the MVA methods is interfaced by the TMVA *Reader*, which is created by the user:

```
TMVA::Reader* reader = new TMVA::Reader( "<options>" );
```

Code Example 21: Instantiating a Reader class object. The only options are the booleans: `V` for verbose, `Color` for coloured output, and `Silent` to suppress all output.

3.3.1 Specifying input variables

The user registers the names of the input variables with the Reader. They are required to be the same (and in the same order) as the names used for training (this requirement is not actually mandatory, but enforced to ensure the consistency between training and application). Together with the name is given the address of a local variable, which carries the updated input values during the event loop.

Macro	Description
`likelihoodrefs.C`	Plots the reference PDFs of all input variables for the projective likelihood method and compares it to original distributions obtained from the training sample.
`network.C`	Draws the TMVA-MLP architecture including weights after training (does not work for the other ANNs).
`annconvergencetest.C`	Plots the MLP error-function convergence versus the training epoch for training and test events (does not work for the other ANNs).
`BDT.C(i)`	Draws the `i`th decision tree of the trained forest (default is `i=1`). The second argument is the weight file that contains the full architecture of the forest (default is `weights/TMVAClassification_BDT.weights.xml`).
`BDTControlPlots.C`	Plots distributions of boost weights throughout forest, boost weights versus decision tree, error fraction, number of nodes before and after pruning and the coefficient α.
`mvarefs.C`	Plots the PDFs used to compute the probability response for a classifier, and compares it to the original distributions.
`PlotFoams.C`	Draws the signal and background foams created by the method PDE-Foam.
`rulevis.C`	Plots the relative importance of rules and linear terms. The 1D plots show the accumulated importance per input variable. The 2D scatter plots show the same but correlated between the input variables. These plots help to identify regions in the parameter space that are important for the model.

Table 4: List of ROOT macros representing results for **specific MVA methods**. The macros require that these methods have been included in the training. All macros take as first argument the name of the ROOT file containing the histograms (default is `TMVA.root`).

```
Int_t   localDescreteVar;
Float_t localFloatingVar, locaSum, localVar3;

reader->AddVariable( "<YourDescreteVar>",                 &localDescreteVar );
reader->AddVariable( "log(<YourFloatingVar>)",            &localFloatingVar );
reader->AddVariable( "SumLabel := <YourVar1>+<YourVar2>", &locaSum          );
reader->AddVariable( "<YourVar3>",                        &localVar3        );
```

Code Example 22: Declaration of the variables and references used as input to the methods (cf. Code Example 11). The order and naming of the variables must be consistent with the ones used for the training. The local variables are updated during the event loop, and through the references their values are known to the MVA methods. The variable type must be either `float` or `int` (`double` is not supported).

3.3.2 Booking MVA methods

The selected MVA methods are booked with the Reader using the weight files from the preceding training job:

```
reader->BookMVA( "<YourMethodName>", "<path/JobName_MethodName.weights.xml>" );
```

Code Example 23: Booking a multivariate method. The first argument is a user defined name to distinguish between methods (it does not need to be the same name as for training, although this could be a useful choice). The true type of the method and its full configuration are read from the weight file specified in the second argument. The default structure of the weight file names is: `path/`⟨`JobName`⟩`_`⟨`MethodName`⟩`.weights.xml`.

3.3.3 Requesting the MVA response

Within the event loop, the response value of a classifier, and – if available – its error, for a given set of input variables computed by the user, are obtained with the commands:

```
localDescreteVar  = treeDescreteVar;       // reference could be implicit
localFloatingVar  = log(treeFloatingVar);
locaSum           = treeVar1 + treeVar2;
localVar3         = treeVar3;              // reference could be implicit

// Classifier response
Double_t mvaValue = reader->EvaluateMVA( "<YourMethodName>" );

// Error on classifier response - must be called after "EvaluateMVA"
// (not available for all methods, returns -1 in that case)
Double_t mvaErr   = reader->GetMVAError();
```

Code Example 24: Updating the local variables for an event, and obtaining the corresponding classifier output and error (if available – see text).

The output of a classifier may then be used for example to put a cut that increases the signal purity of the sample (the achievable purities can be read off the evaluation results obtained during the test phase), or it could enter a subsequent maximum-likelihood fit, or similar. The error reflects the uncertainty, which may be statistical, in the output value as obtained from the training information.

For regression, multi-target response is already supported in TMVA, so that the retrieval command reads:

```
// Regression response for one target
Double_t regValue = (reader->EvaluateRegression( "<YourMethodName>" ))[0];
```

Code Example 25: Obtaining the regression output (after updating the local variables for an event – see above). For mult-target regression, the corresponding vector entries are filled.

The output of a regression method could be directly used for example as energy estimate for a calorimeter cluster as a function of the cell energies.

The rectangular cut classifier is special since it returns a binary answer for a given set of input variables and cuts. The user must specify the desired signal efficiency to define the working point according to which the Reader will choose the cuts:

```
Bool_t passed = reader->EvaluateMVA( "Cuts", signalEfficiency );
```

Code Example 26: For the cut classifier, the second parameter gives the desired signal efficiency according to which the cuts are chosen. The return value is 1 for passed and 0 for retained. See Footnote 19 on page 58 for information on how to determine the optimal working point for known signal and background abundance.

Instead of the classifier response values, one may also retrieve the ratio (12) from the Reader, which, if properly normalised to the expected signal fraction in the sample, corresponds to a probability. The corresponding command reads:

```
Double_t pSig = reader->GetProba( "<YourClassifierName>", sigFrac );
```

Code Example 27: Requesting the event's signal probability from a classifier. The signal fraction is the parameter f_S in Eq. (12).

Similarly, the *Rarity* (13) of a classifier is retrieved by the command

```
Double_t rarity = reader->GetRarity( "<YourClassifierName>" );
```

Code Example 28: Requesting the event's Rarity from a classifier.

3.4 An alternative to the Reader: standalone C++ response classes

To simplify the portability of the trained MVA response to any application the TMVA methods generate after the training lightweight standalone C++ response classes including in the source

```
// load the generated response class into macro and compile it (ROOT)
// or include it into any C++ executable
gROOT->LoadMacro( "TMVAClassification_Fisher.class.C++" ); // usage in ROOT

// define the names of the input variables (same as for training)
std::vector<std::string> inputVars;
inputVars.push_back( "<YourVar1>" );
inputVars.push_back( "log(<YourVar2>)" );
inputVars.push_back( "<YourVar3>+<YourVar4>" );

// create a class object for the Fisher response
IClassifierReader* fisherResponse = new ReadFisher( inputVars );

// the user's event loop ...
std::vector<double> inputVec( 3 );
for (...) {
   // compute the input variables for the event
   inputVec[0] = treeVar1;
   inputVec[1] = TMath::Log(treeVar2);
   inputVec[2] = treeVar3 + treeVar4;

   // get the Fisher response
   double fiOut = fisherResponse->GetMvaValue( inputVec );
   // ... use fiOut
}
```

Code Example 29: Using a standalone C++ class for the classifier response in an application (here of the Fisher discriminant). See also the example code in `TMVA/macros/ClassApplication.C` (Sourceforge.net).

code the content of the weight files.[11] These classes do not depend on ROOT, neither on any other non-standard library. The names of the classes are constructed out of `Read+"MethodName"`, and they inherit from the interface class `IClassifierReader`, which is written into the same C++ file. These standalone classes are *presently only available for classification*.

An example application (ROOT script here, not representative for a C++ standalone application) for a Fisher classifier is given in Code-Example 29. The example is also available in the macro `TMVA/macros/ClassApplication.C` (Sourceforge.net). These classes are C++ representations of the information stored in the weight files. Any change in the training parameters will generate a

[11] At present, the class making functionality has been implemented for all MVA methods with the exception of cut optimisation, PDE-RS, PDE-Foam and k-NN. While for the former classifier the cuts can be easily implemented into the user application, and do not require an extra class, the implementation of a response class for PDE-RS or k-NN requires a copy of the entire analysis code, which we have not attempted so far. We also point out that the use of the standalone C++ class for BDT is not practical due to the colossal size of the generated code.

new class, which must be updated in the corresponding application.[12]

For a given test event, the MVA response returned by the standalone class is identical to the one returned by the Reader. Nevertheless, *we emphasise that the recommended approach to apply the training results is via the Reader*.

[12] We are aware that requiring recompilation constitutes a significant shortcoming. we consider to upgrade these classes to reading the XML weight files, which entails significant complications if the independence of any external library shall be conserved.

4 Data Preprocessing

It is possible to preprocess the discriminating input variables or the training events prior to presenting them to a multivariate method. Preprocessing can be useful to reduce correlations among the variables, to transform their shapes into more appropriate forms, or to accelerate the response time of a method (event sorting).

The preprocessing is completely transparent to the MVA methods. Any preprocessing performed for the training is automatically performed in the application through the Reader class. All the required information is stored in the weight files of the MVA method. Most preprocessing methods discussed below are only available for classification. An exception is the normalisation transformation, which exists for both classification and regression.

4.1 Transforming input variables

Currently four preprocessing transformations are implemented in TMVA:

- variable normalisation;
- decorrelation via the square-root of the covariance matrix ;
- decorrelation via a principal component decomposition;
- transformation of the variables into Gaussian distributions ("Gaussianisation").

With the exception of normalisation, which exists for both classification and regression, the other preprocessing methods are currently only available for classification.

Technically, any transformation of the input variables is performed "on the fly" when the event is requested from the central `DataSet` class. The preprocessing is hence fully transparent to the MVA methods. Any preprocessing performed for the training is automatically also performed in the application through the Reader class. All the required information is stored in the weight files of the MVA method. Each MVA method carries a variable transformation type together with a pointer to the object of its transformation class which is owned by the `DataSet`. If no preprocessing is requested, an identity transform is applied. The `DataSet` registers the requested transformations and takes care not to recreate an identical transformation object (if requested) during the training phase. Hence if two MVA methods wish to apply the same transformation, a single object is shared between them. Each method writes *its* transformation into its weight file once the training has converged. For testing and application of an MVA method, the transformation is read from the weight file and a corresponding transformation object is created. Here each method owns its transformation so that no sharing of potentially different transformation objects occurs (they may have been obtained with different training data and/or under different conditions). A schematic view of the variable transformation interface used in TMVA is drawn in Fig. 10.

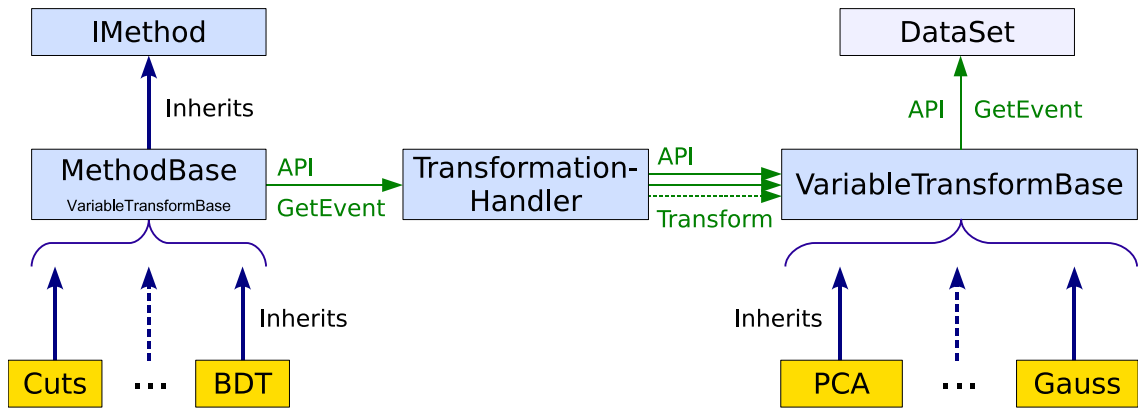

Figure 10: Schematic view of the variable transformation interface implemented in TMVA. Each concrete MVA method derives from `MethodBase` (interfaced by `IMethod`), which holds a protected member object of type `TransformationHandler`. In this object a list of objects derived from `VariableTransformBase` which are the implementations of the particular variable transformations available in TMVA is stored. The construction of the concrete variable transformation objects proceeds in `MethodBase` according to the transformation methods requested in the option string. The events used by the MVA methods for training, testing and final classification (or regression) analysis are read via an API of the `TransformationHandler` class, which itself reads the events from the `DataSet` and applies subsequently all initialised transformations. The `DataSet` fills the current values into the reserved event addresses (the event content may either stem from the training or testing trees, or is set by the user application via the `Reader` for the final classification/regression analysis). The `TransformationHandler` class ensures the proper transformation of all events seen by the MVA methods.

4.1.1 Variable normalisation

Minimum and maximum values for all input variables are determined from the training events and used to linearly scale the input variables to lie within $[-1, 1]$. Such a transformation is useful to allow direct comparisons between MVA weights applied to the variables. Large absolute weights may indicate strong separation power. Normalisation may also render minimisation processes, such as the adjustment of neural network weights, more effective.

4.1.2 Variable decorrelation

A drawback of, for example, the projective likelihood classifier (see Sec. 8.2) is that it ignores correlations among the discriminating input variables. Because in most realistic use cases this is not an accurate conjecture it leads to performance loss. Also other classifiers, such as rectangular cuts or decision trees, and even multidimensional likelihood approaches underperform in presence of variable correlations.

Linear correlations, measured in the training sample, can be taken into account in a straightforward manner through computing the square-root of the covariance matrix. The square-root of a matrix C is the matrix C' that multiplied with itself yields C: $C = (C')^2$. TMVA computes the square-root

matrix by means of diagonalising the (symmetric) covariance matrix

$$D = S^T C S \quad \Rightarrow \quad C' = S\sqrt{D}S^T\,, \tag{14}$$

where D is a diagonal matrix, and where the matrix S is symmetric. The linear decorrelation of the input variables is then obtained by multiplying the initial variable tuple $\mathbf{x}$ by the inverse of the square-root matrix

$$\mathbf{x} \mapsto (C')^{-1}\mathbf{x}\,. \tag{15}$$

The transformations are performed separately for signal and background events because their correlation patterns are usually different.[13] The decorrelation is complete only for linearly correlated and Gaussian distributed variables. In real-world use cases this is not often the case, so that sometimes only little additional information can be recovered by the decorrelation procedure. For highly nonlinear problems the performance may even become worse with linear decorrelation. Nonlinear methods without prior variable decorrelation should be used in such cases.

4.1.3 Principal component decomposition

Principal component decomposition or principal component analysis (PCA) as presently applied in TMVA is not very different from the above linear decorrelation. In common words, PCA is a linear transformation that rotates a sample of data points such that the maximum variability is visible. It thus identifies the most important gradients. In the PCA-transformed coordinate system, the largest variance by any projection of the data comes to lie on the first coordinate (denoted the *first principal component*), the second largest variance on the second coordinate, and so on. PCA can thus be used to reduce the dimensionality of a problem (initially given by the number of input variables) by removing dimensions with insignificant variance. This corresponds to keeping lower-order principal components and ignoring higher-order ones. This latter step however goes beyond straight variable transformation as performed in the preprocessing steps discussed here (it rather represents itself a full classification method). Hence all principal components are retained here.

The tuples $\mathbf{x}_U^{\rm PC}(i) = (x_{U,1}^{\rm PC}(i), \ldots, x_{U,n_{\rm var}}^{\rm PC}(i))$ of principal components of a tuple of input variables $\mathbf{x}(i) = (x_1(i), \ldots, x_{n_{\rm var}}(i))$, measured for the event i for signal ($U = S$) and background ($U = B$), are obtained by the transformation

$$x_{U,k}^{\rm PC}(i) = \sum_{\ell=1}^{n_{\rm var}} (x_{U,\ell}(i) - \overline{x}_{U,\ell})\, v_{U,\ell}^{(k)}\,, \qquad \forall k = 1, n_{\rm var}\,. \tag{16}$$

The tuples $\overline{\mathbf{x}}_U$ and $\mathbf{v}_U^{(k)}$ are the sample means and eigenvectors, respectively. They are computed by the ROOT class `TPrincipal`. The matrix of eigenvectors $V_U = (\mathbf{v}_U^{(1)}, \ldots, \mathbf{v}_U^{(n_{\rm var})})$ obeys the relation

$$C_U \cdot V_U = D_U \cdot V_U\,, \tag{17}$$

[13] Different transformations for signal and background events are only useful for methods that explicitly distinguish signal and background hypotheses. This is the case for the likelihood and PDE-RS classifiers. For all other methods the user must choose which transformation to use.

where C is the covariance matrix of the sample U, and D_U is the tuple of eigenvalues. As for the preprocessing described in Sec. 4.1.2, the transformation (16) eliminates linear correlations for Gaussian variables.

4.1.4 Gaussian transformation of variables ("Gaussianisation")

The decorrelation methods described above require linearly correlated and Gaussian distributed input variables. In real-life HEP applications this is however rarely the case. One may hence transform the variables prior to their decorrelation such that their distributions become Gaussian. The corresponding transformation function is conveniently separated into two steps: first, transform a variable into a uniform distribution using its cumulative distribution function[14] obtained from the training data (this transformation is identical to the "Rarity" introduced in Sec. 3.1.13 on page 28); second, use the inverse error function to transform the uniform distribution into the desired Gaussian shape with zero mean and unity width. As for the other transformations, one needs to choose which class of events is to be transformed and hence, for classification (Gaussianisation is not available for regression), it is only possible to transform signal *or* background into proper Gaussian distributions (except for classifiers testing explicitly both hypotheses such as likelihood methods). Hence a discriminant input variable x with probability density function $\hat{x}$ is transformed as follows

$$x \mapsto \sqrt{2} \cdot \mathrm{erf}^{-1}\left(2 \cdot \int_{-\infty}^{x} \hat{x}(x')\, dx' - 1 \right) . \tag{18}$$

A subsequent decorrelation of the transformed variable tuple sees Gaussian distributions, but most likely non-linear correlations as a consequence of the transformation (18). The distributions obtained after the decorrelation may thus not be Gaussian anymore. It has been suggested that iterating Gaussianisation and decorrelation more than once may improve the performance of likelihood methods (see next section).

4.1.5 Booking and chaining transformations

Variable transformations to be applied prior to the MVA training (and application) can be defined independently for each MVA method with the booking option `VarTransform=<type>`, where `<type>` denotes the desired transformation (or chain of transformations). The available transformation types are normalisation, decorrelation, principal component analysis and Gaussianisation, which are labelled by `Norm`, `Deco`, `PCA`, `Gauss`, respectively, or by the short-hand notation `N`, `D`, `P`, `G`.

Transformations can be chained allowing the consecutive application of all defined transformations to the variables for each event. For example, the above Gaussianisation and decorrelation sequence would be programmed by `VarTransform=G,D`, or even `VarTransform=G,D,G,D` in case of two iterations.

[14] The cumulative distribution function $F(x)$ of the variable x is given by the integral $F(x) = \int_{-\infty}^{x} \hat{x}(x')\, dx'$, where $\hat{x}$ is the probability density function of x.

By default, the transformations are computed with the use of all training events. It is possible to specify the use of a specific class only (e.g., `Signal`, `Background`, `Regression`) by attaching `_<class name>` to the user option – where `<class name>` has to be replaced by the actual class name (e.g., `Signal`) – which defines the transformation (e.g., `VarTransform=G_Signal`). A complex transformation option might hence look like `VarTransform=D,G_Signal,N`. The attachment `_AllClasses` is equivalent to the default, where events from all classes are used.

4.2 Binary search trees

When frequent iterations over the training sample need to be performed, it is helpful to sort the sample before using it. Event sorting in *binary trees* is employed by the MVA methods rectangular cut optimisation, PDE-RS and k-NN. While the former two classifiers rely on the simplest possible binary tree implementation, k-NN uses on a better performing *kd-tree* (cf. Ref. [12]).

Efficiently searching for and counting events that lie inside a multidimensional volume spanned by the discriminating input variables is accomplished with the use of a binary tree search algorithm [13].[15] It is realised in the class `BinarySearchTree`, which inherits from `BinaryTree`, and which is also employed to grow decision trees (cf. Sec. 8.12). The amount of computing time needed to sort N events into the tree is [14] $\propto \sum_{i=1}^{N} \ln_2(i) = \ln_2(N!) \simeq N \ln_2 N$. Finding the events within the tree which lie in a given volume is done by comparing the bounds of the volume with the coordinates of the events in the tree. Searching the tree once requires a CPU time that is $\propto \ln_2 N$, compared to $\propto N^{n_{\rm var}}$ without prior event sorting.

5 Probability Density Functions – the *PDF* Class

Several methods and functionalities in TMVA require the estimation of probability densities (PDE) of one or more correlated variables. One may distinguish three conceptually different approaches to PDEs: (*i*) parametric approximation, where the training data are fitted with a user-defined parametric function, (*ii*) nonparametric approximation, where the data are fitted piecewise using simple standard functions, such as a polynomial or a Gaussian, and (*iii*) nearest-neighbour estimation, where the average of the training data in the vicinity of a test event determines the PDF.

All multidimensional PDEs used in TMVA are based on nearest-neighbour estimation with however quite varying implementations. They are described in Secs. 8.3, 8.4 and 8.5.

[15] The following is extracted from Ref. [14] for a two-dimensional range search example. Consider a random sequence of signal events $e_i(x_1, x_2)$, $i = 1, 2, \ldots$, which are to be stored in a binary tree. The first event in the sequence becomes by definition the topmost node of the tree. The second event $e_2(x_1, x_2)$ shall have a larger x_1-coordinate than the first event, therefore a new node is created for it and the node is attached to the first node as the right child (if the x_1-coordinate had been smaller, the node would have become the left child). Event e_3 shall have a larger x_1-coordinate than event e_1, it therefore should be attached to the right branch below e_1. Since e_2 is already placed at that position, now the x_2-coordinates of e_2 and e_3 are compared, and, since e_3 has a larger x_2, e_3 becomes the right child of the node with event e_2. The tree is sequentially filled by taking every event and, while descending the tree, comparing its x_1 and x_2 coordinates with the events already in place. Whether x_1 or x_2 are used for the comparison depends on the level within the tree. On the first level, x_1 is used, on the second level x_2, on the third again x_1 and so on.

One-dimensional PDFs in TMVA are estimated by means of nonparametric approximation, because parametric functions cannot be generalised to a-priori unknown problems. The training data can be in form of binned histograms, or unbinned data points (or "quasi-unbinned" data, i.e., histograms with very narrow bins). In the first case, the bin centres are interpolated with polynomial spline curves, while in the latter case one attributes a kernel function to each event such that the PDF is represented by the sum over all kernels. Beside a faithful representation of the training data, it is important that statistical fluctuations are smoothed out as much as possible without destroying significant information. In practise, where the true PDFs are unknown, a compromise determines which information is regarded significant and which is not. Likelihood methods crucially depend on a good-quality PDF representation. Since the PDFs are strongly problem dependent, the default configuration settings in TMVA will almost never be optimal. The user is therefore advised to scrutinise the agreement between training data and PDFs via the available plotting macros, and to optimise the settings.

In TMVA, all PDFs are derived from the `PDF` class, which is instantiated via the command (usually hidden in the MVA methods for normal TMVA usage):

```
PDF* pdf = new PDF( "<options>", "Suffix", defaultPDF );
pdf->BuildPDF( SourceHistogram );
double p = pdf->GetVal( x );
```

Code Example 30: Creating and using a PDF class object. The first argument is the configuration options string. Individual options are separated by a ':'. The second optional argument is the suffix appended to the options used in the option string. The suffix is added to the option names given in the Option Table 3 in order to distinguish variables and types. The third (optional) object is a PDF from which default option settings are read. The histogram specified in the second line is a TH1* object from which the PDF is built. The third line shows how to retrieve the PDF value at a given test value `'x'`.

Its configuration options are given in Option Table 3.

5.1 Nonparametric PDF fitting using spline functions

Polynomial splines of various orders are fitted to one-dimensional (1D) binned histograms according to the following procedure.

1. The number of bins of the `TH1` object representing the distribution of the input variable is driven by the options `NAvEvtPerBin` or `Nbins` (cf. Option Table 3). Setting `Nbins` enforces a fixed number of bins, while `NAvEvtPerBin` defines an average number of entries required per bin. The upper and lower bounds of the histogram coincide with the limits found in the data (or they are $[-1, 1]$ if the input variables are normalised).

2. The histogram is smoothed adaptively between `MinNSmooth` and `MaxNSmooth` times, using `TH1::Smooth(.)` – an implementation of the 353QH-twice algorithm [20]. The appropriate

Option	Values	Description
PDFInterpol	KDE, Spline0, Spline1, Spline2*, Spline3, Spline5	The method of interpolating the reference histograms: either by using the unbinned kernel density estimator (KDE), or by various degrees of spline functions (note that currently the KDE characteristics cannot be changed individually but apply to all variables that select KDE)
NSmooth	0	Number of smoothing iterations for the input histograms; if set, MinNSmooth and MaxNSmooth are ignored
MinNSmooth	-1	Minimum number of smoothing iteration for the input histograms; for bins with least relative error (see text)
MaxNSmooth	-1	Maximum number of smoothing iteration for the input histograms; for bins with most relative error (see text)
Nbins	0	Number of bins used to build the reference histogram; if set to value > 0, NAvEvtPerBin is ignored
NAvEvtPerBin	50	Average number of events per bin in the reference histogram (see text)
KDEtype	Gauss*	KDE kernel type (currently only Gauss)
KDEiter	Nonadaptive*, Adaptive	Non-adaptive or adaptive number of iterations (see text)
KDEFineFactor	1	Fine-tuning factor for the adaptive KDE
KDEborder	None*, Renorm, Mirror	Method for correcting histogram boundary/border effects
CheckHist	False	Sanity comparison of derived high-binned interpolated PDF histogram versus the original PDF function

Option Table 3: Configuration options for class: *PDF*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. In case a suffix is defined for the PDF, it is added in the end of the option name. For example for PDF with suffix MVAPdf the number of smoothing iterations is given by the option NSmoothMVAPdf (see Option Table 11 on page 62 for a concrete use of the PDF options in a MVA method).

number of smoothing iterations is derived with the aim to preserve statistically significant structures, while smoothing out fluctuations. Bins with the largest (smallest) relative statistical error are maximally (minimally) smoothed. All other bins are smoothed between MaxNSmooth and MinNSmooth times according to a linear function of their relative errors. During the smoothing process a histogram with the suffix NSmooth is created for each variable, where the number of smoothing iterations applied to each bin is stored.

3. The smoothed TH1 object is internally cloned and the requested polynomial splines are used to interpolate adjacent bins. All spline classes are derived from ROOT's TSpline class. Available are: polynomials of degree 0 (the original smoothed histogram is kept), which is used for discrete variables; degree 1 (linear), 2 (quadratic), 3 (cubic) and degree 5. Splines of degree two or more render the PDF continuous and differentiable in all points excluding the interval

borders. In case of a likelihood analysis, this ensures the same property for the likelihood ratio (34). Since cubic (and higher) splines equalise the first and second derivatives at the spline transitions, the resulting curves, although mathematically smooth, can wiggle in quite unexpected ways. Furthermore, there is no local control of the spline: moving one control point (bin) causes the entire curve to change, not just the part near the control point. To ensure a safe interpolation, quadratic splines are used by default.

4. To speed up the numerical access to the probability densities, the spline functions are stored into a finely binned (10^4 bins) histogram, where adjacent bins are interpolated by a linear function. Only after this step, the PDF is normalised according to Eq. (36).

5.2 Nonparametric PDF parameterisation using kernel density estimators

Another type of nonparametric approximation of a 1D PDF are kernel density estimators (KDE). As opposed to splines, KDEs are obtained from unbinned data. The idea of the approach is to estimate the shape of a PDF by the sum over *smeared* training events. One then finds for a PDF $p(x)$ of a variable x [21]

$$p(x) = \frac{1}{N\,h}\sum_{i=1}^{N} K\left(\frac{x-x_i}{h}\right) = \frac{1}{N}\sum_{i=1}^{N} K_h(x-x_i)\ , \tag{19}$$

where N is the number of training events, $K_h(t) = K(t/h)/h$ is the kernel function, and h is the *bandwidth* of the kernel (also termed the *smoothing parameter*). Currently, only a Gaussian form of K is implemented in TMVA, where the exact form of the kernel function is of minor relevance for the quality of the shape estimation. More important is the choice of the bandwidth.

The KDE smoothing can be applied in either non-adaptive (NA) or adaptive form (A), the choice of which is controlled by the option `KDEiter`. In the non-adaptive case the bandwidth $h_{\rm NA}$ is kept constant for the entire training sample. As optimal bandwidth can be taken the one that minimises the *asymptotic mean integrated square error* (AMISE). For the case of a Gaussian kernel function this leads to [21]

$$h_{\rm NA} = \left(\frac{4}{3}\right)^{1/5} \sigma_x N^{-1/5}\,, \tag{20}$$

where σ_x is the RMS of the variable x.

The so-called *sample point adaptive* method uses as input the result of the non-adaptive KDE, but also takes into account the local event density. The adaptive bandwidth $h_{\rm A}$ then becomes a function of $p(x)$ [21]

$$h_{\rm A}(x) = \frac{h_{\rm NA}}{\sqrt{p(x)}}\,. \tag{21}$$

The adaptive approach improves the shape estimation in regions with low event density. However, in regions with high event density it can give rise to "over-smoothing" of fine structures such as narrow peaks. The degree of smoothing can be tuned by multiplying the bandwidth $h_{\rm A}(x)$ with the user-specified factor `KDEFineFactor`.

For practical reasons, the KDE implementation in TMVA differs somewhat from the procedure described above. Instead of unbinned training data a finely-binned histogram is used as input, which allows to significantly speed up the algorithm. The calculation of the optimal bandwidth h_{NA} is performed in the dedicated class `KDEKernel`. If the algorithm is run in the adaptive mode, the non-adaptive step is also performed because its output feeds the computation of $h_{\mathrm{A}}(x)$ for the adaptive part. Subsequently, a smoothed high-binned histogram estimating the PDF shape is created by looping over the bins of the input histogram and summing up the corresponding kernel functions, using h_{NA} ($h_{\mathrm{A}}(x)$) in case of the non-adaptive (adaptive) mode. This output histogram is returned to the `PDF` class.

Both the non-adaptive and the adaptive methods can suffer from the so-called *boundary problem* at the histogram boundaries. It occurs for instance if the original distribution is non-zero below a physical boundary value and zero above. This property cannot be reproduced by the KDE procedure. In general, the stronger the discontinuity the more acute is the boundary problem. TMVA provides three options under the term `KDEborder` that allow to treat boundary problems.

- `KDEborder=None`
  No boundary treatment is performed. The consequence is that close to the boundary the KDE result will be inaccurate: below the boundary it will underestimate the PDF while it will not drop to zero above. In TMVA the PDF resulting from KDE is a (finely-binned) histogram, with bounds equal to the minimum and the maximum values of the input distribution. Hence, the boundary value will be at the edge of the PDF (histogram), and a drop of the PDF due to the proximity of the boundary can be observed (while the artificial enhancement beyond the boundary will fall outside of the histogram). In other words, for training events that are close to the boundary some fraction of the probability "flows" outside the histogram (*probability leakage*). As a consequence, the integral of the kernel function inside the histogram borders is smaller than one.

- `KDEborder=Renorm`
  The probability leakage is compensated by renormalising the kernel function such that the integral inside the histogram borders is equal to one.

- `KDEborder=Mirror`
  The original distribution is "mirrored" around the boundary. The same procedure is applied to the mirrored events and each of them is smeared by a kernel function and its contribution *inside the histogram's (PDF) boundaries* is added to the PDF. The mirror copy compensates the probability leakage completely.

6 Optimisation and Fitting

Several MVA methods (notably cut optimisation and FDA) require general purpose parameter fitting to optimise the value of an estimator. For example, an estimator could be the sum of the deviations of classifier outputs from 1 for signal events and 0 for background events, and the parameters are adjusted so that this sum is as small as possible. Since the various fitting problems

Option	Values	Description
`FitMethod`	`MC, MINUIT, GA, SA`	Fitter method
`Converger`	`None*, MINUIT`	Converger which can be combined with MC or GA (currently only used for FDA) to improve finding local minima

Option Table 4: Configuration options for the choice of a fitter. The abbreviations stand for Monte Carlo sampling, Minuit, Genetic Algorithm, Simulated Annealing. By setting a Converger (only Minuit is currently available) combined use of Monte Carlo sampling and Minuit, and of Genetic Algorithm and Minuit is possible. The `FitMethod` option can be used in any MVA method that requires fitting. The option `Converger` is currently only implemented in FDA. The default fitter depends on the MVA method. The fitters and their specific options are described below.

call for dedicated solutions, TMVA has a fitter base class, used by the MVA methods, from which all concrete fitters inherit. The consequence of this is that the user can choose whatever fitter is deemed suitable and can configure it through the option string of the MVA method. At present, four fitters are implemented and described below: Monte Carlo sampling, Minuit minimisation, a Genetic Algorithm, Simulated Annealing. They are selected via the configuration option of the corresponding MVA method for which the fitter is invoked (see Option Table 4). Combinations of MC and GA with Minuit are available for the FDA method by setting the `Converger` option, as described in Option Table 16.

6.1 Monte Carlo sampling

The simplest and most straightforward, albeit inefficient fitting method is to randomly sample the fit parameters and to choose the set of parameters that optimises the estimator. The priors used for the sampling are uniform or Gaussian within the parameter limits. The specific configuration options for the MC sampling are given in Option Table 5.

For fitting problems with few local minima out of which one is a global minimum the performance can be enhanced by setting the parameter `Sigma` to a positive value. The newly generated parameters are then not any more independent of the parameters of the previous samples. The random generator will throw random values according to a Gaussian probability density with the mean given by the currently known best value for that particular parameter and the width in units of the interval size given by the option `Sigma`. Points which are created out of the parameter's interval are mapped back into the interval.

6.2 Minuit minimisation

Minuit is the standard multivariate minimisation package used in HEP [17]. Its purpose is to find the minimum of a multi-parameter estimator function and to analyse the shape of the function around the minimum (error analysis). The principal application of the TMVA fitters is simple

Option	Array	Default	Predefined Values	Description
`SampleSize`	–	`100000`	–	Number of Monte Carlo events in toy sample
`Sigma`	–	`-1`	–	If > 0: new points are generated according to Gauss around best value and with Sigma in units of interval length
`Seed`	–	`100`	–	Seed for the random generator (0 takes random seeds)

Option Table 5: Configuration options reference for fitting method: *Monte Carlo sampling (MC)*.

minimisation, while the shape of the minimum is irrelevant in most cases. The use of Minuit is therefore not necessarily the most efficient solution, but because it is a very robust tool we have included it here. Minuit searches the solution along the direction of the gradient until a minimum or an boundary is reached (MIGRAD algorithm). It does not attempt to find the global minimum but is satisfied with local minima. If during the error analysis with MINOS, the minimum smaller values than the local minimum might be obtained. In particular, the use of MINOS may as a side effect of an improved error analysis uncover a convergence in a local minimum, in which case MIGRAD minimisation is invoked once again. If multiple local and/or global solutions exist, it might be preferable to use any of the other fitters which are specifically designed for this type of problem.

The configuration options for Minuit are given in Option Table 6.

6.3 Genetic Algorithm

A Genetic Algorithm is a technique to find approximate solutions to optimisation or search problems. The problem is modelled by a group (*population*) of abstract representations (*genomes*) of possible solutions (*individuals*). By applying means similar to processes found in biological evolution the individuals of the population should evolve towards an optimal solution of the problem. Processes which are usually modelled in evolutionary algorithms — of which Genetic Algorithms are a subtype — are inheritance, mutation and "sexual recombination" (also termed *crossover*).

Apart from the abstract representation of the solution domain, a *fitness* function must be defined. Its purpose is the evaluation of the goodness of an individual. The fitness function is problem dependent. It either returns a value representing the individual's goodness or it compares two individuals and indicates which of them performs better.

The Genetic Algorithm proceeds as follows:

- *Initialisation*: A starting population is created. Its size depends on the problem to be solved. Each individual belonging to the population is created by randomly setting the parameters of

Option	Array	Default	Predefined Values	Description
`ErrorLevel`	–	`1`	–	TMinuit: error level: 0.5=logL fit, 1=chi-squared fit
`PrintLevel`	–	`-1`	–	TMinuit: output level: -1=least, 0, +1=all garbage
`FitStrategy`	–	`2`	–	TMinuit: fit strategy: 2=best
`PrintWarnings`	–	`False`	–	TMinuit: suppress warnings
`UseImprove`	–	`True`	–	TMinuit: use IMPROVE
`UseMinos`	–	`True`	–	TMinuit: use MINOS
`SetBatch`	–	`False`	–	TMinuit: use batch mode
`MaxCalls`	–	`1000`	–	TMinuit: approximate maximum number of function calls
`Tolerance`	–	`0.1`	–	TMinuit: tolerance to the function value at the minimum

Option Table 6: Configuration options reference for fitting method: *Minuit*. More information on the Minuit parameters can be found here: http://root.cern.ch/root/html/TMinuit.html.

the abstract representation (variables), thus producing a point in the solution domain of the initial problem.

- *Evaluation*: Each individual is evaluated using the fitness function.
- *Selection*: Individuals are kept or discarded as a function of their fitness. Several selection procedures are possible. The simplest one is to separate out the worst performing fraction of the population. Another possibility is to decide on the individual's survival by assigning probabilities that depend on the individual's performance compared to the others.
- *Reproduction*: The surviving individuals are copied, mutated and crossed-over until the initial population size is reached again.
- *Termination*: The evaluation, selection and reproduction steps are repeated until a maximum number of cycles is reached or an individual satisfies a maximum-fitness criterion. The best individual is selected and taken as solution to the problem.

The TMVA Genetic Algorithm provides controls that are set through configuration options (cf. Table 7). The parameter `PopSize` determines the number of individuals created at each generation of the Genetic Algorithm. At the initialisation, all parameters of all individuals are chosen randomly. The individuals are evaluated in terms of their fitness, and each individual giving an improvement is immediately stored.

Individuals with a good fitness are selected to engender the next generation. The new individuals are created by crossover and mutated afterwards. Mutation changes some values of some parameters of some individuals randomly following a Gaussian distribution function. The width of the Gaussian

Option	Array	Default	Predefined Values	Description
`PopSize`	–	`300`	–	Population size for GA
`Steps`	–	`40`	–	Number of steps for convergence
`Cycles`	–	`3`	–	Independent cycles of GA fitting
`SC_steps`	–	`10`	–	Spread control, steps
`SC_rate`	–	`5`	–	Spread control, rate: factor is changed depending on the rate
`SC_factor`	–	`0.95`	–	Spread control, factor
`ConvCrit`	–	`0.001`	–	Convergence criteria
`SaveBestGen`	–	`1`	–	Saves the best n results from each generation. They are included in the last cycle
`SaveBestCycle`	–	`10`	–	Saves the best n results from each cycle. They are included in the last cycle. The value should be set to at least 1.0
`Trim`	–	`False`	–	Trim the population to PopSize after assessing the fitness of each individual
`Seed`	–	`100`	–	Set seed of random generator (0 gives random seeds)

Option Table 7: Configuration options reference for fitting method: *Genetic Algorithm (GA).*

can be altered by the parameter `SC\_factor`. The current width is multiplied by this factor when within the last `SC\_steps` generations more than `SC\_rate` improvements have been obtained. If there were `SC\_rate` improvements the width remains unchanged. Were there, on the other hand, less than `SC\_rate` improvements, the width is divided by `SC\_factor`. This allows to influence the speed of searching through the solution domain.

The cycle of evaluating the fitness of the individuals of a generation and producing a new generation is repeated until the improvement of the fitness within the last `Steps` has been less than `ConvCrit`. The minimisation is then considered to have converged. The whole cycle from initialisation over fitness evaluation, selection, reproduction and determining the improvement is repeated `Cycles` times, before the Genetic Algorithm has finished.

Guidelines for optimising the GA

`PopSize` is the most important value for enhancing the quality of the results. This value is by default set to 300, but can be increased to 1000 or more only limited by the resources available. The calculation time of the GA should increase with O(`PopSize`).

`Steps` is set by default to 40. This value can be increased moderately to about 60. Time consumption increases non linearly but at least with O(`Steps`).

`Cycles` is set by default to 3. In this case, the GA is called three times independently from each other. With `SaveBestCycle` and `SaveBestGen` it is possible to set the number of best results which shall be stored each cycle of even each generation. These stored results are reused in the last cycle. That way the last cycle is not independent from the others, but incorporates their best results. The number of cycles can be increased moderately to about 10. The time consumption of GA rises with about O(`Cycles`).

6.4 Simulated Annealing

Simulated Annealing also aims at solving a minimisation problem with several discrete or continuous, local or global minima. The algorithm is inspired by the process of of annealing which occur in condensed matter physics. When first heating and then slowly cooling down a metal ("annealing") its atoms move towards a state of lowest energy, while for sudden cooling the atoms tend to freeze in intermediate states higher energy. For infinitesimal annealing activity the system will always converge in its global energy minimum (see, e.g., Ref. [18]). This physical principle can be converted into an algorithm to achieve slow, but correct convergence of an optimisation problem with multiple solutions. Recovery out of local minima is achieved by assigning the probability [19]

$$p(\Delta E) \propto \exp\left(-\frac{\Delta E}{T}\right) , \tag{22}$$

to a perturbation of the parameters leading to a shift ΔE in the energy of the system. The probability of such perturbations to occur decreases with the size of a positive energy coefficient of the perturbation, and increases with the ambient temperature (T).

Guidelines for optimising SA

The TMVA implementation of Simulated Annealing includes various different adaptive adjustments of the perturbation and temperature gradients. The adjustment procedure is chosen by setting `KernelTemp` to one of the following values.

- **Increasing Adaptive Approach** (`IncAdaptive`). The algorithm seeks local minima and explores their neighbourhood, while changing the ambient temperature depending on the number of failures in the previous steps. The performance can be improved by increasing the number of iteration steps (`MaxCalls`), or by adjusting the minimal temperature (`MinTemp`). Manual adjustments of the speed of the temperature increase (`TempScale` and `AdaptiveSpeed`) for individual data sets might also improve the performance.

- **Decreasing Adaptive Approach** (`DecAdaptive`). The algorithm calculates the initial temperature (based on the effectiveness of large steps) and defines a multiplier to ensure that the minimal temperature is reached in the requested number of iteration steps. The performance can be improved by adjusting the minimal temperature (`MinTemp`) and by increasing number of steps (`MaxCalls`).

Option	Array	Default	Predefined Values	Description
`MaxCalls`	–	`100000`	–	Maximum number of minimisation calls
`InitialTemp`	–	`1e+06`	–	Initial temperature
`MinTemp`	–	`1e-06`	–	Mimimum temperature
`Eps`	–	`1e-10`	–	Epsilon
`TempScale`	–	`1`	–	Temperature scale
`AdaptiveSpeed`	–	`1`	–	Adaptive speed
`TempAdaptiveStep`	–	`0.009875`	–	Step made in each generation temperature adaptive
`UseDefaultScale`	–	`False`	–	Use default temperature scale for temperature minimisation algorithm
`UseDefaultTemp`	–	`False`	–	Use default initial temperature
`KernelTemp`	–	`IncAdaptive`	`IncAdaptive, DecAdaptive, Sqrt, Log, Sin, Homo, Geo`	Temperature minimisation algorithm

Option Table 8: Configuration options reference for fitting method: *Simulated Annealing (SA)*.

- **Other Kernels.** Several other procedures to calculate the temperature change are also implemented. Each of them starts with the maximum temperature (`MaxTemp`) and decreases while changing the temperature according to :

$$\text{Temperature}^{(k)} = \begin{cases} \texttt{Sqrt}: & \frac{\texttt{InitialTemp}}{\sqrt{k+2}} \cdot \texttt{TempScale} \\ \texttt{Log}: & \frac{\texttt{InitialTemp}}{\ln(k+2)} \cdot \texttt{TempScale} \\ \texttt{Homo}: & \frac{\texttt{InitialTemp}}{k+2} \cdot \texttt{TempScale} \\ \texttt{Sin}: & \frac{\sin(k/\texttt{TempScale})+1}{k+1} \cdot \texttt{InitialTemp} + \varepsilon \\ \texttt{Geo}: & \texttt{CurrentTemp} \cdot \texttt{TempScale} \end{cases}$$

  Their performances can be improved by adjusting the initial temperature `InitialTemp` ($= \text{Temperature}^{(k=0)}$), the number of iteration steps (`MaxCalls`), and the multiplier that scales the temperature decrease (`TempScale`).

The configuration options for the Simulated Annealing fitter are given in Option Table 8.

6.5 Combined fitters

For MVA methods such as FDA, where parameters of a discrimination function are adjusted to achieve optimal classification performance (cf. Sec. 8.9), the user can choose to combine Minuit

parameter fitting with Monte Carlo sampling or a Genetic Algorithm. While the strength of Minuit is the speedy detection of a nearby local minimum, it might not find a better global minimum. If several local minima exist Minuit will find different solutions depending on the start values for the fit parameters. When combining Minuit with Monte Carlo sampling or a Genetic Algorithm, Minuit uses starting values generated by these methods. The subsequent fits then converge in local minima. Such a combination is usually more efficient than the uncontrolled sampling used in Monte Carlo techniques. When combined with a Genetic Algorithm the user can benefit from the advantages of both methods: the Genetic Algorithm to roughly locate the global minimum, and Minuit to find an accurate solution for it (for an example see the FDA method).

The configuration options for the combined fit methods are the inclusive sum of all the individual fitter options. It is recommended to use Minuit in batch mode (option `SetBatch`) and without MINOS (option `!UseMinos`) to prevent TMVA from flooding the output with Minuit messages which cannot be turned off, and to speed up the individual fits. It is however important to note that the combination of MIGRAD and MINOS together is less susceptible to getting caught in local minima.

7 Boosting and Bagging

Boosting is a way of enhancing the classification and regression performance (and increasing the stability with respect to statistical fluctuations in the training sample) of typically weak MVA methods by sequentially applying an MVA algorithm to reweighted (*boosted*) versions of the training data and then taking a weighted majority vote of the sequence of MVA algorithms thus produced. It has been introduced to classification techniques in the early '90s [25] and in many cases this simple strategy results in dramatic performance increases.

Although one may argue that *bagging* (cf. Sec. 7.3) is not a genuine boosting algorithm, we include it in the same context and typically when discussing boosting we also refer to bagging. The most commonly boosted methods are decision trees. However, as described in Sec. 9.1 on page 119, any MVA method may be boosted with TMVA. Hence, although the following discussion refers to decision trees, it also applies to other methods. (Note however that "Gradient Boost" is only available for decision trees and only for classification in the present TMVA version).

7.1 Adaptive Boost (AdaBoost)

The most popular boosting algorithm is the so-called *AdaBoost* (adaptive boost) [26]. In a *classification problem*, events that were misclassified during the training of a decision tree are given a higher event weight in the training of the following tree. Starting with the original event weights when training the first decision tree, the subsequent tree is trained using a modified event sample where the weights of previously misclassified events are multiplied by a common *boost weight* α.

The boost weight is derived from the misclassification rate, err, of the previous tree[16],

$$\alpha = \frac{1 - \text{err}}{\text{err}} \,. \tag{23}$$

The weights of the entire event sample are then renormalised such that the sum of weights remains constant.

We define the result of an individual classifier as $h(\mathbf{x})$, with ($\mathbf{x}$ being the tuple of input variables) encoded for signal and background as $h(\mathbf{x}) = +1$ and -1, respectively. The boosted event classification $y_{\text{Boost}}(\mathbf{x})$ is then given by

$$y_{\text{Boost}}(\mathbf{x}) = \frac{1}{N_{\text{collection}}} \cdot \sum_{i}^{N_{\text{collection}}} \ln(\alpha_i) \cdot h_i(\mathbf{x}) \,, \tag{24}$$

where the sum is over all classifiers in the collection. Small (large) values for $y_{\text{Boost}}(\mathbf{x})$ indicate a background-like (signal-like) event. Equation (24) represents the default boosting algorithm. It can be modified via the configuration option string of the MVA method to be boosted (see Option Tables 21 and 22 on pages 106 and 106 for boosted decision trees, and Option Table 24 for general classifier boosting 9.1) if one wants to use an unweighted average of the boosted decision trees or classifiers instead of the weighted one.

For *regression trees*, the AdaBoost algorithm needs to be modified. TMVA uses here the so-called AdaBoost.R2 algorithm [27]. The idea is similar to AdaBoost albeit with a redefined loss per event to account for the the deviation of the estimated target value from the true one. Moreover, as there are no longer correctly and wrongly classified events, all events need to be reweighted depending on their individual *loss*, which – for event k – is given by

$$\textit{Linear}: \qquad L(k) = \frac{|y(k) - \hat{y}(k)|}{\max\limits_{\text{events } k'} (|y(k') - \hat{y}(k'|)} \,, \tag{25}$$

$$\textit{Square}: \qquad L(k) = \left[\frac{|y(k) - \hat{y}(k)|}{\max\limits_{\text{events } k'} (|y(k') - \hat{y}(k'|)} \right]^2 \,, \tag{26}$$

$$\textit{Exponential}: \quad L(k) = 1 - \exp\left[- \frac{|y(k) - \hat{y}(k)|}{\max\limits_{\text{events } k'} (|y(k') - \hat{y}(k'|)} \right] \,. \tag{27}$$

The average loss of the classifier $y^{(i)}$ over the whole training sample, $\langle L \rangle^{(i)} = \sum_{\text{events } k'} w(k') L^{(i)}(k')$, can be considered to be the analogon to the error fraction in classification. Given $\langle L \rangle$, one computes the quantity $\beta_{(i)} = \langle L \rangle^{(i)} / (1 - \langle L \rangle^{(i)})$, which is used in the boosting of the events, and for the combination of the regression methods belonging to the boosted collection. The boosting weight, $w^{(i+1)}(k)$, for event k and boost step $i+1$ thus reads

$$w^{(i+1)}(k) = w^{(i)}(k) \cdot \beta_{(i)}^{1 - L^{(i)}(k)} \,. \tag{28}$$

[16] By construction, the error rate is err ≤ 0.5 as the same training events used to classify the output nodes of the previous tree are used for the calculation of the error rate.

The sum of the event weights is again renormalised to reproduce the original overall number of events. The final regressor, y_{Boost}, uses the weighted median, $\tilde{y}_{(i)}$, where (i) is chosen so that it is the minimal (i) that satisfies the inequality

$$\sum_{\substack{t \in \text{sorted collection} \\ t \leq i}} \ln \frac{1}{\beta_{(t)}} \geq \frac{1}{2} \sum_{t}^{N_{\text{collection}}} \ln \frac{1}{\beta_{(t)}} \tag{29}$$

7.2 Gradient Boost

The idea of function estimation through boosting can be understood by considering a simple additive expansion approach. The function $F(\mathbf{x})$ under consideration is assumed to be a weighted sum of parametrised base functions $f(x; a_m)$, so-called "weak learners". From a technical point of view any TMVA classifier could act as a weak learner in this approach, but decision trees benefit most from boosting and are currently the only classifier that implements GradientBoost (a generalisation may be included in future releases). Thus each base function in this expansion corresponds to a decision tree

$$F(\mathbf{x}; P) = \sum_{m=0}^{M} \beta_m f(x; a_m); \qquad P \in \{\beta_m; a_m\}_0^M \ . \tag{30}$$

The boosting procedure is now employed to adjust the parameters P such that the deviation between the model response $F(\mathbf{x})$ and the true value y obtained from the training sample is minimised. The deviation is measured by the so-called *loss-function* $L(F, y)$, a popular choice being squared error loss $L(F, y) = (F(\mathbf{x}) - y)^2$. It can be shown that the loss function fully determines the boosting procedure.

The most popular boosting method, AdaBoost, is based on exponential loss, $L(F, y) = e^{-F(\mathbf{x})y}$, which leads to the well known reweighting algorithm described in Sec. 7.1. Exponential loss has the shortcoming that it lacks robustness in presence of outliers or mislabelled data points. The performance of AdaBoost therefore degrades in noisy settings.

The *GradientBoost* algorithm attempts to cure this weakness by allowing for other, potentially more robust, loss functions without giving up on the good out-of-the-box performance of AdaBoost. The current TMVA implementation of GradientBoost uses the binomial log-likelihood loss

$$L(F, y) = \ln \left(1 + e^{-2F(\mathbf{x})y}\right) \ , \tag{31}$$

for classification. As the boosting algorithm corresponding to this loss function cannot be obtained in a straightforward manner, one has to resort to a steepest-descent approach to do the minimisation. This is done by calculating the current gradient of the loss function and then growing a regression tree whose leaf values are adjusted to match the mean value of the gradient in each region defined by the tree structure. Iterating this procedure yields the desired set of decision trees which minimises the loss function. Note that GradientBoost can be adapted to any loss function as long as the calculation of the gradient is feasible.

Giving good results already for small trees (5–10 leaf nodes), GradientBoost is typically less susceptible to overtraining. Its robustness can be enhanced by reducing the learning rate of the algorithm through the `Shrinkage` parameter (cf. Option Table 21 on page 106), which controls the weight of the individual trees. A small shrinkage (0.1–0.3) demands more trees to be grown but can significantly improve the accuracy of the prediction in difficult settings.

In certain settings GradientBoost may also benefit from the introduction of a bagging-like resampling procedure using random subsamples of the training events for growing the trees. This is called *stochastic gradient boosting* and can be enabled by selecting the `UseBaggedGrad` option. The sample fraction used in each iteration can be controlled through the parameter `GradBaggingFraction`, where typically the best results are obtained for values between 0.5 and 0.8.

For regression tasks, GradientBoost employs the Huber loss function [45], which features the robustness of least-absolute-deviation loss for error distributions with broad tails, while maintaining the high efficiency of least-squares loss for normally distributed errors. For moderately broad tails, it should surpass both least-squares and least-absolute-deviation loss.

$$L(F,y) = \begin{cases} \frac{1}{2}\left(y - F(\mathbf{x})\right)^2 & |y-F| \leq \delta\,, \\ \delta(|y-F| - \delta/2) & |y-F| > \delta\,. \end{cases} \tag{32}$$

All GradientBoost options for classification described above can be also applied for regression tasks, but tests have shown that stochastic gradient boosting may not perform very well for regression problems.

7.3 Bagging

The term *Bagging* denotes a resampling technique where a classifier is repeatedly trained using resampled training events such that the combined classifier represents an average of the individual classifiers. A priori, bagging does not aim at enhancing a weak classifier in the way adaptive or gradient boosting does, and is thus not a "boosting" algorithm in a strict sense. Instead it effectively smears over statistical representations of the training data and is hence suited to stabilise the response of a classifier. In this context it is often accompanied also by a significant performance increase compared to the individual classifier.

Resampling includes the possibility of replacement, which means that the same event is allowed to be (randomly) picked several times from the parent sample. This is equivalent to regarding the training sample as being a representation of the probability density distribution of the parent sample: indeed, if one draws an event out of the parent sample, it is more likely to draw an event from a region of phase-space that has a high probability density, as the original data sample will have more events in that region. If a selected event is kept in the original sample (that is when the same event can be selected several times), the parent sample remains unchanged so that the randomly extracted samples will have the same parent distribution, albeit statistically fluctuated. Training several classifiers with different resampled training data, and combining them into a collection, results in an averaged classifier that, just as for boosting, is more stable with respect to statistical fluctuations in the training sample.

Technically, resampling is implemented by applying random Poisson weights to each event of the parent sample.

8 The TMVA Methods

All TMVA classification and regression methods (in most cases, a method serves both analysis goals) inherit from `MethodBase`, which implements basic functionality like the interpretation of common configuration options, the interaction with the training and test data sets, I/O operations and common performance evaluation calculus. The functionality each MVA method is required to implement is defined in the abstract interface `IMethod`.[17] Each MVA method provides a function that creates a rank object (of type `Ranking`), which is an ordered list of the input variables prioritised according to criteria specific to that method. Also provided are brief method-specific help notes (option `Help`, switched off by default) with information on the adequate usage of the method and performance optimisation in case of unsatisfying results.

If the option `CreateMVAPdfs` is set TMVA creates signal and background PDFs from the corresponding MVA response distributions using the training sample (cf. Sec. 3.1.7). The binning and smoothing properties of the underlying histograms can be customised via controls implemented in the `PDF` class (cf. Sec. 5 and Option Table 3 on page 43). The options specific to `MethodBase` are listed in Option Table 9. They can be accessed by all MVA methods.

The following sections describe the methods implemented in TMVA. For each method we proceed according to the following scheme: (*i*) a brief introduction, (*ii*) the description of the booking options required to configure the method, (*iii*) a description of the the method and TMVA implementation specifications for classification and – where available – for regression, (*iv*) the properties of the variable ranking, and (*v*) a few comments on performance, favourable (and disfavoured) use cases, and comparisons with other methods.

8.1 Rectangular cut optimisation

The simplest and most common classifier for selecting signal events from a mixed sample of signal and background events is the application of an ensemble of rectangular cuts on discriminating variables. Unlike all other classifiers in TMVA, the cut classifier only returns a binary response (signal *or* background).[18] The optimisation of cuts performed by TMVA maximises the background rejection at given signal efficiency, and scans over the full range of the latter quantity. Dedicated analysis optimisation for which, e.g., the signal *significance* is maximised requires the expected signal and background yields to be known before applying the cuts. This is not the case for a

[17] Two constructors are implemented for each method: one that creates the method for a first time for training with a configuration ("option") string among the arguments, and another that recreates a method from an existing weight file. The use of the first constructor is demonstrated in the example macros `TMVAClassification.C` and `TMVARegression.C`, while the second one is employed by the Reader in `TMVAClassificationApplication.C` and `TMVARegressionApplication.C`. Other functions implemented by each methods are: `Train` (called for training), `Write/ReadWeightsToStream` (I/O of specific training results), `WriteMonitoringHistosToFile` (additional specific information for monitoring purposes) and `CreateRanking` (variable ranking).

[18] Note that cut optimisation is not a *multivariate* analyser method but a sequence of univariate ones, because no combination of the variables is achieved. Neither does a cut on one variable depend on the value of another variable (like it is the case for Decision Trees), nor can a, say, background-like value of one variable in a signal event be counterweighed by signal-like values of the other variables (like it is the case for the likelihood method).

Option	Array	Default	Predefined Values	Description
`V`	–	`False`	–	Verbose output (short form of VerbosityLevel below - overrides the latter one)
`VerbosityLevel`	–	`Default`	`Default, Debug, Verbose, Info, Warning, Error, Fatal`	Verbosity level
`VarTransform`	–	`None`	–	List of variable transformations performed before training, e.g., D_Background,P_Signal,G,N_AllClasses for: Decorrelation, PCA-transformation, Gaussianisation, Normalisation, each for the given class of events ('AllClasses' denotes all events of all classes, if no class indication is given, 'All' is assumed)
`H`	–	`False`	–	Print method-specific help message
`CreateMVAPdfs`	–	`False`	–	Create PDFs for classifier outputs (signal and background)
`IgnoreNegWeightsInTraining`		`False`	–	Events with negative weights are ignored in the training (but are included for testing and performance evaluation)

Option Table 9: Configuration options that are common for all classifiers (but which can be controlled individually for each classifier). Values given are defaults. If predefined categories exist, the default category is marked by a *. The lower options in the table control the PDF fitting of the classifiers (required, e.g., for the Rarity calculation).

multi-purpose discrimination and hence not used by TMVA. However, the cut ensemble leading to maximum significance corresponds to a particular working point on the efficiency curve, and can hence be easily derived after the cut optimisation scan has converged.[19]

TMVA cut optimisation is performed with the use of multivariate parameter fitters interfaced by the class `FitterBase` (cf. Sec. 6). Any fitter implementation can be used, where however because of the peculiar, non-unique solution space only Monte Carlo sampling, Genetic Algorithm, and

[19] Assuming a large enough number of events so that Gaussian statistics is applicable, the significance for a signal is given by $\mathcal{S} = \varepsilon_{\mathcal{S}}\mathcal{N}_{\mathcal{S}}/\sqrt{\varepsilon_{\mathcal{S}}\mathcal{N}_{\mathcal{S}} + \varepsilon_{\mathcal{B}}(\varepsilon_{\mathcal{S}})\mathcal{N}_{\mathcal{S}}}$, where $\varepsilon_{S(B)}$ and $N_{S(B)}$ are the signal and background efficiencies for a cut ensemble and the event yields before applying the cuts, respectively. The background efficiency ε_B is expressed as a function of ε_S using the TMVA evaluation curve obtained form the test data sample. The maximum significance is then found at the root of the derivative

$$\frac{d\mathcal{S}}{d\varepsilon_S} = N_S \frac{2\varepsilon_B(\varepsilon_S)N_B + \varepsilon_S\left(N_S - \frac{d\varepsilon_B(\varepsilon_S)}{d\varepsilon_S}N_B\right)}{2\left(\varepsilon_S N_S + \varepsilon_B(\varepsilon_S)N_B\right)^{3/2}} = 0\,, \tag{33}$$

which depends on the problem.

Option	Array	Default	Predefined Values	Description
`FitMethod`	–	`GA`	`GA, SA, MC, MCEvents, MINUIT, EventScan`	Minimisation Method (GA, SA, and MC are the primary methods to be used; the others have been introduced for testing purposes and are deprecated)
`EffMethod`	–	`EffSel`	`EffSel, EffPDF`	Selection Method
`CutRangeMin`	Yes	`-1`	–	Minimum of allowed cut range (set per variable)
`CutRangeMax`	Yes	`-1`	–	Maximum of allowed cut range (set per variable)
`VarProp`	Yes	`NotEnforced`	`NotEnforced, FMax, FMin, FSmart, FVerySmart`	Categorisation of cuts

Option Table 10: Configuration options reference for MVA method: *Cuts*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured.

Simulated Annealing show satisfying results. Attempts to use Minuit (SIMPLEX or MIGRAD) have not shown satisfactory results, with frequently failing fits.

The training events are sorted in *binary trees* prior to the optimisation, which significantly reduces the computing time required to determine the number of events passing a given cut ensemble (cf. Sec. 4.2).

8.1.1 Booking options

The rectangular cut optimisation is booked through the Factory via the command:

```
factory->BookMethod( Types::kCuts, "Cuts", "<options>" );
```

Code Example 31: Booking of the cut optimisation classifier: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The configuration options for the various cut optimisation techniques are given in Option Table 10.

8.1.2 Description and implementation

The cut optimisation analysis proceeds by first building binary search trees for signal and background. For each variable, statistical properties like mean, root-mean-squared (RMS), variable ranges are computed to guide the search for optimal cuts. Cut optimisation requires an estimator that quantifies the goodness of a given cut ensemble. Maximising this estimator minimises (maximises) the background efficiency, ε_B (background rejection, $r_B = 1 - \varepsilon_B$) for each signal efficiency ε_S.

All optimisation methods (fitters) act on the assumption that one minimum and one maximum requirement on each variable is sufficient to optimally discriminate signal from background (i.e., the signal is clustered). If this is not the case, the variables must be transformed prior to the cut optimisation to make them compliant with this assumption.

For a given cut ensemble the signal and background efficiencies are derived by counting the training events that pass the cuts and dividing the numbers found by the original sample sizes. The resulting efficiencies are therefore rational numbers that may exhibit visible discontinuities when the number of training events is small and an efficiency is either very small or very large. Another way to compute efficiencies is to parameterise the probability density functions of all input variables and thus to achieve continuous efficiencies for any cut value. Note however that this method expects the input variables to be uncorrelated! Non-vanishing correlations would lead to incorrect efficiency estimates and hence to underperforming cuts. The two methods are chosen with the option `EffMethod` set to `EffSel` and `EffPDF`, respectively.

Monte Carlo sampling

Each generated cut sample (cf. Sec. 6.1) corresponds to a point in the (ε_S, r_B) plane. The ε_S dimension is (finely) binned and a cut sample is retained if its r_B value is larger than the value already contained in that bin. This way a reasonably smooth efficiency curve can be obtained if the number of input variables is not too large (the required number of MC samples grows with powers of $2n_{\text{var}}$).

Prior information on the variable distributions can be used to reduce the number of cuts that need to be sampled. For example, if a discriminating variable follows Gaussian distributions for signal and background, with equal width but a larger mean value for the background distribution, there is no useful minimum requirement (other than $-\infty$) so that a single maximum requirement is sufficient for this variable. To instruct TMVA to remove obsolete requirements, the option `VarProp[i]` must be used, where `[i]` indicates the counter of the variable (following the order in which they have been registered with the Factory, beginning with 0) must be set to either `FMax` or `FMin`. TMVA is capable of automatically detecting which of the requirements should be removed. Use the option `VarProp[i]=FSmart` (where again `[i]` must be replaced by the appropriate variable counter, beginning with 0). Note that in many realistic use cases the mean values between signal and background of a variable are indeed distinct, but the background can have large tails. In such a case, the removal of a requirement is inappropriate, and would lead to underperforming cuts.

Genetic Algorithm

Genetic Algorithm (cf. Sec. 6.3) is a technique to find approximate solutions to optimisation or search problems. Apart from the abstract representation of the solution domain, a *fitness* function must be defined. In cut optimisation, the fitness of a rectangular cut is given by good background rejection combined with high signal efficiency.

At the initialization step, all parameters of all individuals (cut ensembles) are chosen randomly. The individuals are evaluated in terms of their background rejection and signal efficiency. Each cut ensemble giving an improvement in the background rejection for a specific signal efficiency bin is immediately stored. Each individual's fitness is assessed, where the fitness is largely determined by the difference of the best found background rejection for a particular bin of signal efficiency and the value produced by the current individual. The same individual that has at one generation a very good fitness will have only average fitness at the following generation. This forces the algorithm to focus on the region where the potential of improvement is the highest. Individuals with a good fitness are selected to produce the next generation. The new individuals are created by crossover and mutated afterwards. Mutation changes some values of some parameters of some individuals randomly following a Gaussian distribution function, etc. This process can be controlled with the parameters listed in Option Table 7, page 49.

Simulated Annealing

Cut optimisation using Simulated Annealing works similarly as for the Genetic Algorithm and achieves comparable performance. The same fitness function is used to estimator the goodness of a given cut ensemble. Details on the algorithm and the configuration options can be found in Sec. 6.4 on page 50.

8.1.3 Variable ranking

The present implementation of Cuts does not provide a ranking of the input variables.

8.1.4 Performance

The Genetic Algorithm currently provides the best cut optimisation convergence. However, it is found that with rising number of discriminating input variables the goodness of the solution found (and hence the smoothness of the background-rejections versus signal efficiency plot) deteriorates quickly. Rectangular cut optimisation should therefore be reduced to the variables that have the largest discriminating power.

If variables with excellent signal from background separation exist, applying cuts can be quite competitive with more involved classifiers. Cuts are known to underperform in presence of strong nonlinear correlations and/or if several weakly discriminating variables are used. In the latter case,

Option	Array	Default	Predefined Values	Description
`TransformOutput`	–	`False`	–	Transform likelihood output by inverse sigmoid function

Option Table 11: Configuration options reference for MVA method: *Likelihood*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured.

a true multivariate combination of the information will be rewarding.

8.2 Projective likelihood estimator (PDE approach)

The method of maximum likelihood consists of building a model out of probability density functions (PDF) that reproduces the input variables for signal and background. For a given event, the likelihood for being of signal type is obtained by multiplying the signal probability densities of all input variables, which are assumed to be independent, and normalising this by the sum of the signal and background likelihoods. Because correlations among the variables are ignored, this PDE approach is also called "naive Bayes estimator", unlike the full multidimensional PDE approaches such as PDE-range search, PDE-foam and k-nearest-neighbour discussed in the subsequent sections, which approximate the true Bayes limit.

8.2.1 Booking options

The likelihood classifier is booked via the command:

```
factory->BookMethod( Types::kLikelihood, "Likelihood", "<options>" );
```

Code Example 32: Booking of the (projective) likelihood classifier: the first argument is the predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The likelihood configuration options are given in Option Table 11.

8.2.2 Description and implementation

The likelihood ratio $y_{\mathcal{L}}(i)$ for event i is defined by

$$y_{\mathcal{L}}(i) = \frac{\mathcal{L}_S(i)}{\mathcal{L}_S(i) + \mathcal{L}_B(i)}\,, \tag{34}$$

where

$$\mathcal{L}_{S(B)}(i) = \prod_{k=1}^{n_{\rm var}} p_{S(B),k}(x_k(i))\,, \tag{35}$$

and where $p_{S(B),k}$ is the signal (background) PDF for the kth input variable x_k. The PDFs are normalised

$$\int_{-\infty}^{+\infty} p_{S(B),k}(x_k)dx_k = 1\,, \quad \forall k. \tag{36}$$

It can be shown that in absence of model inaccuracies (such as correlations between input variables not removed by the de-correlation procedure, or an inaccurate probability density model), the ratio (34) provides optimal signal from background separation for the given set of input variables.

Since the parametric form of the PDFs is generally unknown, the PDF shapes are empirically approximated from the training data by nonparametric functions, which can be chosen individually for each variable and are either polynomial splines of various degrees fitted to histograms or unbinned kernel density estimators (KDE), as discussed in Sec. (5).

A certain number of primary validations are performed during the PDF creation, the results of which are printed to standard output. Among these are the computation of a χ^2 estimator between all nonzero bins of the original histogram and its PDF, and a comparison of the number of outliers (in sigmas) found in the original histogram with respect to the (smoothed) PDF shape, with the statistically expected one. The fidelity of the PDF estimate can be also inspected visually by executing the macro `likelihoodrefs.C` (cf. Table 4).

Transforming the likelihood output

If a data-mining problem offers a large number of input variables, or variables with excellent separation power, the likelihood response $y_{\mathcal{L}}$ is often strongly peaked at 0 (background) and 1 (signal). Such a response is inconvenient for the use in subsequent analysis steps. TMVA therefore allows to transform the likelihood output by an inverse sigmoid function that zooms into the peaks

$$y_{\mathcal{L}}(i) \longrightarrow y'_{\mathcal{L}}(i) = -\tau^{-1}\ln\big(y_{\mathcal{L}}^{-1} - 1\big)\ , \tag{37}$$

where $\tau = 15$ is used. Note that $y'_{\mathcal{L}}(i)$ is no longer contained within $[0,1]$ (see Fig. 11). The transformation (37) is enabled (disabled) with the booking option `TransformOutput=True(False)`.

8.2.3 Variable ranking

The present likelihood implementation does not provide a ranking of the input variables.

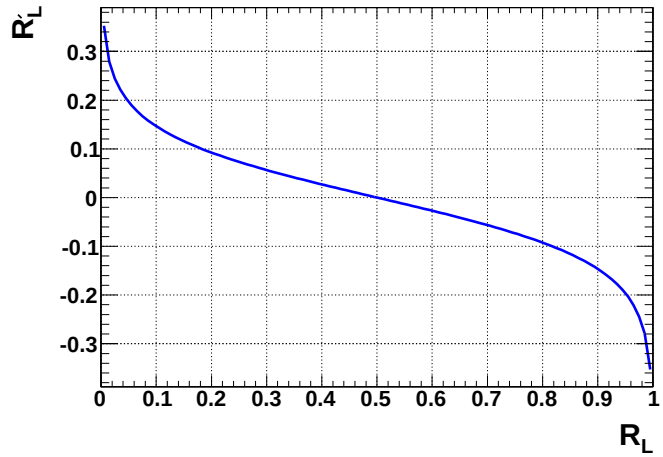

Figure 11: Transformation (37) of the likelihood output.

8.2.4 Performance

Both the training and the application of the likelihood classifier are very fast operations that are suitable for large data sets.

The performance of the classifier relies on the accuracy of the likelihood model. Because high fidelity PDF estimates are mandatory, sufficient training statistics is required to populate the tails of the distributions. The neglect of correlations between input variables in the model (35), often leads to a diminution of the discrimination performance. While linear Gaussian correlations can be rotated away (see Sec. 4.1), such an ideal situation is rarely given. Positive correlations lead to peaks at both $y_{\mathcal{L}} \to 0, 1$. Correlations can be reduced by categorising the data samples and building an independent likelihood classifier for each event category. Such categories could be geometrical regions in the detector, kinematic properties, etc. In spite of this, realistic applications with a large number of input variables are often plagued by irreducible correlations, so that projective likelihood approaches like the one discussed here are under-performing. This finding led to the development of the many alternative classifiers that exist in statistical theory today.

8.3 Multidimensional likelihood estimator (PDE range-search approach)

This is a generalization of the projective likelihood classifier described in Sec. 8.2 to n_{var} dimensions, where n_{var} is the number of input variables used. If the multidimensional PDF for signal and background (or regression data) were known, this classifier would exploit the full information contained in the input variables, and would hence be optimal. In practice however, huge training samples are necessary to sufficiently populate the multidimensional phase space.[20] Kernel estimation methods may be used to approximate the shape of the PDF for finite training statistics.

A simple probability density estimator denoted *PDE range search*, or *PDE-RS*, has been suggested

[20] Due to correlations between the input variables, only a sub-space of the full phase space may be populated.

in Ref. [14]. The PDE for a given test event (discriminant) is obtained by counting the (normalised) number of training events that occur in the "vicinity" of the test event. The classification of the test event may then be conducted on the basis of the majority of the nearest training events. The n_{var}-dimensional volume that encloses the "vicinity" is user-defined and can be adaptive. A search method based on sorted binary trees is used to reduce the computing time for the range search. To enhance the sensitivity within the volume, kernel functions are used to weight the reference events according to their distance from the test event. PDE-RS is a variant of the k-nearest neighbour classifier described in Sec. 8.5.

8.3.1 Booking options

The PDE-RS classifier is booked via the command:

```
factory->BookMethod( Types::kPDERS, "PDERS", "<options>" );
```

Code Example 33: Booking of PDE-RS: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The configuration options for the PDE-RS classifier are given in Option Table 12.

8.3.2 Description and implementation

Classification

To classify an event as being either of signal or of background type, a *local* estimate of the probability density of it belonging to either class is computed. The method of PDE-RS provides such an estimate by defining a volume (V) around the test event (i), and by counting the number of signal ($n_S(i,V)$) and background events ($n_B(i,V)$) obtained from the training sample in that volume. The ratio

$$y_{\text{PDE-RS}}(i,V) = \frac{1}{1+r(i,V)}\,, \tag{38}$$

is taken as the estimate, where $r(i,V) = (n_B(i,V)/N_B)\cdot(N_S/n_S(i,V))$, and $N_{S(B)}$ is the total number of signal (background) events in the training sample. The estimator $y_{\text{PDE-RS}}(i,V)$ peaks at 1 (0) for signal (background) events. The counting method averages over the PDF within V, and hence ignores the available shape information inside (and outside) that volume.

Binary tree search

Efficiently searching for and counting the events that lie inside the volume is accomplished with the use of a n_{var}-variable binary tree search algorithm [13] (cf. Sec. 4.2).

Option	Array	Default	Predefined Values	Description
`VolumeRangeMode`	–	`Adaptive`	`Unscaled, MinMax, RMS, Adaptive, kNN`	Method to determine volume size
`KernelEstimator`	–	`Box`	`Box, Sphere, Teepee, Gauss, Sinc3, Sinc5, Sinc7, Sinc9, Sinc11, Lanczos2, Lanczos3, Lanczos5, Lanczos8, Trim`	Kernel estimation function
`DeltaFrac`	–	`3`	–	nEventsMin/Max for minmax and rms volume range
`NEventsMin`	–	`100`	–	nEventsMin for adaptive volume range
`NEventsMax`	–	`200`	–	nEventsMax for adaptive volume range
`MaxVIterations`	–	`150`	–	MaxVIterations for adaptive volume range
`InitialScale`	–	`0.99`	–	InitialScale for adaptive volume range
`GaussSigma`	–	`0.1`	–	Width (wrt volume size) of Gaussian kernel estimator
`NormTree`	–	`False`	–	Normalize binary search tree

Option Table 12: Configuration options reference for MVA method: *PDE-RS*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured.

Choosing a volume

The TMVA implementation of PDE-RS optionally provides four different volume definitions selected via the configuration option `VolumeRangeMode`.

- `Unscaled`
  The simplest volume definition consisting of a rigid box of size `DeltaFrac`, in units of the variables. This method was the one originally used by the developers of PDE-RS [14].

- `MinMax`
  The volume is defined in each dimension (i.e., input variable) with respect to the full range of values found for that dimension in the training sample. The fraction of this volume used for the range search is defined by the option `DeltaFrac`.

- `RMS`
  The volume is defined in each dimension with respect to the RMS of that dimension (input

variable), estimated from the training sample. The fraction of this volume used for the range search is defined by the option `DeltaFrac`.

- `Adaptive`
  A volume is defined in each dimension with respect to the RMS of that dimension, estimated from the training sample. The overall scale of the volume is adjusted individually for each test event such that the total number of events confined in the volume lies within a user-defined range (options `NEventsMin/Max`). The adjustment is performed by the class `RootFinder`, which is a C++ implementation of Brent's algorithm (translated from the CERNLIB function RZERO). The maximum initial volume (fraction of the RMS) and the maximum number of iterations for the root finding is set by the options `InitialScale` and `MaxVIterations`, respectively. The requirement to collect a certain number of events in the volume automatically leads to small volume sizes in strongly populated phase space regions, and enlarged volumes in areas where the population is scarce.

Although the adaptive volume adjustment is more flexible and should perform better, it significantly increases the computing time of the PDE-RS discriminant. If found too slow, one can reduce the number of necessary iterations by choosing a larger `NEventsMin/Max` interval.

Event weighting with kernel functions

One of the shortcomings of the original PDE-RS implementation is its sensitivity to the exact location of the sampling volume boundaries: an infinitesimal change in the boundary placement can include or exclude a training event, thus changing $r(i, V)$ by a finite amount.[21] In addition, the shape information within the volume is ignored.

Kernel functions mitigate these problems by weighting each event within the volume as a function of its distance to the test event. The farer it is away, the smaller is its weight. The following kernel functions are implemented in TMVA, and can be selected with the option `KernelEstimator`.

- `Box`
  Corresponds to the original rectangular volume element without application of event weights.
- `Sphere`
  A hyper-elliptic volume element is used without application of event weights. The hyper-ellipsoid corresponds to a sphere of constant fraction in the `MinMax` or `RMS` metrics. The size of the sphere can be chosen adaptive, just as for the rectangular volume.
- `Teepee`
  The simplest linear interpolation that eliminates the discontinuity problem of the box. The training events are given a weight that decreases linearly with their distance from the centre of the volume (the position of the test event). In other words: these events are convolved with the triangle or tent function, becoming a sort of teepee in multi-dimensions.

[21] Such an introduction of artefacts by having sharp boundaries in the sampled space is an example of Gibbs's phenomenon, and is commonly referred to as *ringing* or *aliasing*.

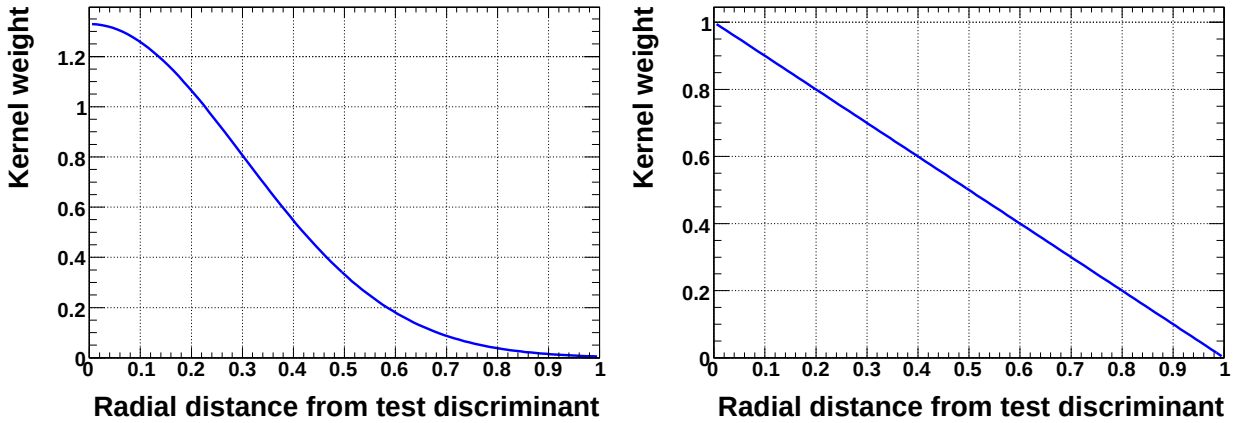

Figure 12: Kernel functions (left: Gaussian, right: Teepee) used to weight the events that are found inside the reference volume of a test event.

- `Trim`
  Modified Teepee given by the function $(1 - \hat{x}^3)^3$, where $\hat{x}$ is the the normalised distance between test event and reference.

- `Gauss`
  The simplest well behaved convolution kernel. The width of the Gaussian (fraction of the volume size) can be set by the option `GaussSigma`.

Other kernels implemented for test purposes are "Sinc" and "Lanczos" functions $\propto \sin x/x$ of different (symmetric) orders. They exhibit strong peaks at zero and oscillating tails. The Gaussian and Teepee kernel functions are shown in Fig. 12.

Regression

Regression with PDE-RS proceeds similar to classification. The difference lies in the replacement of Eq. (38) by the average target value of all events belonging to the volume V defined by event i (the *test* event)

$$y_{\text{PDE-RS}}(i, V) = \langle t(i, V) \rangle = \frac{\sum_{j \in V} w_j t_j f(\text{dis}(i, j))}{\sum_{j \in V} w_j f(\text{dis}(i, j))} \,, \tag{39}$$

where the sum is over all training events in V, w_j and t_j are the weight and target value of event j in V, $\text{dis}(i, j)$ is a measure of the distance between events i and j, and $f(\dots)$ is a kernel function.

8.3.3 Variable ranking

The present implementation of PDE-RS does not provide a ranking of the input variables.

8.3.4 Performance

As opposed to many of the more sophisticated data-mining approaches, which tend to present the user with a "black box", PDE-RS is simple enough that the algorithm can be easily traced and tuned by hand. PDE-RS can yield competitive performance if the number of input variables is not too large and the statistics of the training sample is ample. In particular, it naturally deals with complex nonlinear variable correlations, the reproduction of which may, for example, require involved neural network architectures.

PDE-RS is a slowly responding classifier. Only the training, i.e., the fabrication of the binary tree is fast, which is usually not the critical part. The necessity to store the entire binary tree in memory to avoid accessing virtual memory limits the number of training events that can effectively be used to model the multidimensional PDF. This is not the case for the other classifiers implemented in TMVA (with some exception for Boosted Decision Trees).

8.4 Likelihood estimator using self-adapting phase-space binning (PDE-Foam)

The PDE-Foam method [16] is an extension of PDE-RS, which divides the multi-dimensional phase space in a finite number of hyper-rectangles (cells) of constant event density. This "foam" of cells is filled with averaged probability density information sampled from the training data. For a given number of cells, the binning algorithm adjusts the size and position of the cells inside the multi-dimensional phase space based on a binary split algorithm that minimises the variance of the event density in the cell. The binned event density information of the final foam is stored in cells, organised in a binary tree, to allow a fast and memory-efficient storage and retrieval of the event density information necessary for classification or regression. The implementation of PDE-Foam is based on the Monte-Carlo integration package TFoam [15] included in ROOT.

In classification mode PDE-Foam forms bins of similar density of signal and background events or the ratio of signal to background. In regression mode the algorithm determines cells with small varying regression targets. In the following, we use the term density (ρ) for the event density in case of classification or for the target variable density in case of regression.

8.4.1 Booking options

The PDE-Foam classifier is booked via the command:

```
factory->BookMethod( Types::kPDEFoam, "PDEFoam", "<options>" );
```

Code Example 34: Booking of PDE-Foam: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

Option	Array	Default	Predefined Values	Description
SigBgSeparate	–	False	–	Separate foams for signal and background
TailCut	–	0.001	–	Fraction of outlier events that are excluded from the foam in each dimension
VolFrac	–	0.0333333	–	Size of sampling box, used for density calculation during foam build-up (maximum value: 1.0 is equivalent to volume of entire foam)
nActiveCells	–	500	–	Maximum number of active cells to be created by the foam
nSampl	–	2000	–	Number of generated MC events per cell
nBin	–	5	–	Number of bins in edge histograms
Compress	–	True	–	Compress XML file
MultiTargetRegression	–	False	–	Do regression with multiple targets
CutNmin	–	True	–	Requirement for minimal number of events in cell
Nmin	–	100	–	Number of events in cell required to split cell
Kernel	–	None	None, Gauss, LinNeighbors	Kernel type used
TargetSelection	–	Mean	Mean, Mpv	Target selection method

Option Table 13: Configuration options reference for MVA method: *PDE-Foam*. The options in Option Table 9 on page 58 can also be configured.

The configuration options for the PDE-Foam method are summarised in Option Table 13.

Table 5 gives an overview of supported combinations of configuration options.

8.4.2 Description and implementation of the foam algorithm

Foams for an arbitrary event sample are formed as follows.

1. *Setup of binary search trees.* A binary search tree is created and filled with the d-dimensional event tuples form the training sample as for the PDE-RS method (cf. Sec. 8.3).

2. *Initialisation phase.* A foam is created, which at first consists of one d-dimensional hyperrectangle (base cell). The coordinate system of the foam is normalised such that the base cell extends from 0 to 1 in each dimension. The coordinates of the events in the corresponding training tree are linearly transformed into the coordinate system of the foam.

Option	Classification		Regression	
	Separated foams	Single foam	Mono target	Multi target
`SigBgSeparate`	True	False	–	–
`MultiTargetRegression`	–	–	False	True
`CutNmin`	•	•	•	•
`Kernel`	•	•	•	•
`TargetSelection`	○	○	○	•
`TailCut`	•	•	•	•

Table 5: Availability of options for the two classification and two regression modes of PDE-Foam. Supported options are marked by a '•', while disregarded ones are marked by a '○'.

3. *Growing phase.* A binary splitting algorithm iteratively splits cells of the foam along axis-parallel hyperplanes until the maximum number of active cells (set by `nActiveCells`) is reached. The splitting algorithm minimises the relative variance of the density $\sigma_\rho/\langle\rho\rangle$ across each cell (cf. Ref. [15]). For each cell `nSampl` random points uniformly distributed over the cell volume are generated. For each of these points a small box centred around this point is defined. The box has a size of `VolFrac` times the size of the base cell in each dimension. The density is estimated as the number of events contained in this box divided by the volume of the box.[22] The densities obtained for all sampled points in the cell are projected onto the d axes of the cell and the projected values are filled in histograms with `nBin` bins. The next cell to be split and the corresponding division edge (bin) for the split are selected as the ones that have the largest relative variance. The two new daughter cells are marked as 'active' cells and the old mother cell is marked as 'inactive'. A detailed description of the splitting algorithm can be found in Ref. [15]. The geometry of the final foam reflects the distribution of the training samples: phase-space regions where the density is constant are combined in large cells, while in regions with large density gradients many small cells are created. Figure 13(a) displays a foam obtained from a two-dimensional Gaussian-distributed training sample.

4. *Filling phase.* Each active cell is filled with values that classify the event distribution within this cell and which permit the calculation of the classification or regression discriminator.

5. *Evaluation phase.* The estimator for a given event is evaluated based on the information stored in the foam cells. The corresponding foam cell, in which the event variables (d-dimensional vector) of a given event is contained, is determined with a binary search algorithm.[23]

The initial trees which contain the training events and which are needed to evaluate the densities for the foam build-up, are discarded after the training phase. The memory consumption for the foam is 160 bytes per foam cell plus an overhead of 1.4 kbytes for the PDE-Foam object on a 64 bit architecture. Note that in the foam all cells created during the growing phase are stored within a binary tree structure. Cells which have been split are marked as inactive and remain empty. To

[22] In case of regression this is the average target value computed according to Eq. (39), page 68.

[23] For an event that falls outside the foam boundaries, the cell with the smallest Cartesian distance to the event is chosen.

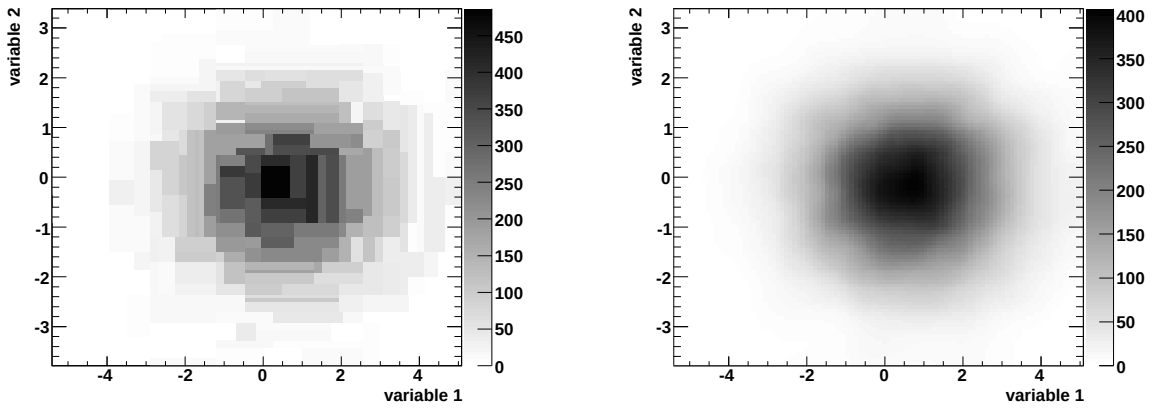

(a) foam projection without kernel

(b) foam projection with Gaussian kernel

Figure 13: Projections of a two-dimensional foam with 500 cells for a Gaussian distribution on a two-dimensional histogram. The foam was created with 5000 events from the input tree. (a) shows the reconstructed distribution without kernel weighting and (b) shows the distribution weighted with a Gaussian kernel. The grey shades indicate the event density of the drawn cell. For more details about the projection function see the description on page 78.

reduce memory consumption, the cell geometry is not stored with the cell, but rather obtained recursively from the information about the division edge of the corresponding mother cell. This way only two short integer numbers per cell represent the information of the entire foam geometry: the division coordinate and the bin number of the division edge.

PDE-Foam options

- `TailCut` – *boundaries for foam geometry*
  A parameter `TailCut` is used to define the geometry of the base cell(s) such that outliers of the distributions in the training ensemble are not included in the volume of the base cell(s). In a first step, the upper and lower limits for each input variable are determined from the training sample. Upper and a lower bounds are then determined for each input variable, such that on both sides a fraction `TailCut` of all events is excluded. The default value of `TailCut=0.001` ensures a sufficient suppression of outliers for most input distributions. For cases where the input distributions have a fixed range or where they are discontinuous and/or have peaks towards the edges, it can be appropriate to set `TailCut` to 0. Note that for event classification it is guaranteed that the foam has an infinite coverage: events outside the foam volume are assigned to the cell with the smallest Cartesian distance to the event.

- `nActiveCells` – *maximum number of active foam cells*
  In most cases larger `nActiveCells` values result in more accurate foams, but also lead to longer computation time during the foam formation, and require more storage memory. The default value of 500 was found to be a good compromise for most practical applications if the size of the training samples is of the order of 10^4 events. Note that the total number of cells,

`nCells`, is given as `nCells` = `nActiveCells` $\cdot 2 - 1$, since the binary-tree structure of the foam requires all inactive cells to remain part of the foam (see *Growing phase*).

- `VolFrac` – *size of the probe volume for the density sampling of the training data*
  The volume is defined as a d-dimensional box with edge length `VolFrac` times the extension of the base cell in each dimension. The default of 1/30 results in a box with volume $1/30^d$ times the volume of the base cell. Smaller values of `VolFrac` increase the sampling speed, but also make the algorithm more vulnerable to statistical fluctuations in the training sample (overtraining). In cases with many observables (> 5) and small training samples ($< 10^4$), `VolFrac` should be increased for better classification results.

- `nSampl` – *number of samplings per cell and per cell-division step*
  The computation time for the foam formation scales linearly with the number of samplings. The default of 2000 was found to give sufficiently accurate estimates of the density distributions with an acceptable computation time.[24]

- `nBin` – *number of bins for edge histograms*
  The number of bins for the edge histograms used to determine the variance of the sampled density distributions within one cell are set with the parameter `nBin`. The performance in typical applications was found to be rather insensitive to the number of bins. The default value of 5 gives the foam algorithm sufficient flexibility to find the division points, while maintaining sufficient statistical accuracy also in case of small event samples.

- `Nmin` – *Minimal number of events for cell split*
  If the option `CutNmin=T` is set, the foam will only consider cells with a number of events greater than `Nmin` to be split. If no more cells are found during the growing phase for which this condition is fulfilled, the foam will stop splitting cells, even if the target number of cells is not yet reached. This option prevents the foam from adapting to statistical fluctuations in the training samples (overtraining). Note that event weights are not taken into account for evaluating the number of events in the cell.

  In particular for training samples with small event numbers of less than 10^4 events this cut improves the performance. The default value of (`Nmin=100`) was found to give good results in most cases. It should be reduced in applications with very small training samples (less than 200 training events) and with many cells.

- `Kernel` – *cell weighting with kernel functions:*
  A Gaussian Kernel smoothing is applied during the evaluation phase, if the option `Kernel` is set to "`Gauss`". In this case all cells contribute to the calculation of the discriminant for a given event, weighted with their Gaussian distance to the event. The average cell value v (event density in case of separated foams, and the ratio $n_S/(n_S+n_B)$ in case of a single foam) for a given event $\mathbf{x} = (x_1, \ldots, x_{n_{\text{var}}})$ is calculated by

$$v = \frac{\sum_{\text{all cells } i} G(D(i,\mathbf{x}), 0, \texttt{VolFrac}) \cdot v_i}{\sum_{\text{all cells } i} G(D(i,\mathbf{x}), 0, \texttt{VolFrac})}\,, \tag{40}$$

[24] Contrary to the original application where an analytical function is used, increasing the number of samplings does not automatically lead to better accuracy. Since the density is defined by a limited event sample, it may occur that always the same event is found for all sample points.

where v_i is the output value in cell i, $G(x, x_0, \sigma) = \exp(-(x - x_0)^2/2\sigma^2)$, and $D(i, \mathbf{x})$ is the minimal Euclidean distance between $\mathbf{x}$ and any point $\mathbf{y}$ in cell i

$$D(i, \mathbf{x}) = \min_{\mathbf{y} \in \text{cell } i} |\mathbf{x} - \mathbf{y}| \,. \tag{41}$$

The Gaussian kernel avoids discontinuities of the discriminant values at the cell boundaries. In most cases it results in an improved separation power between signal and background. However, the time needed for classification increases due to the larger number of computations performed. A comparison between foams with and without Gaussian kernel can be seen in Fig. 13.

A linear interpolation with adjacent cells in each dimension is applied during the classification phase, if the option `Kernel` is set to "`LinNeighbors`". This results in faster classification than the Gauss weighting of all cells in the foam.

The PDE-Foam algorithm exhibits stable performance with respect to variations in most of the parameter settings. However, optimisation of the parameters is required for small training samples ($< 10^4$ events) in combination with many observables (> 5). In such cases, `VolFrac` should be increased until an acceptable performance is reached. Moreover, in cases where the classification time is not critical, one of the `Kernel` methods should be applied to further improve the classification performance. For large training samples ($> 10^5$) and if the training time is not critical, `nActiveCells` should be increased to improve the classification performance.

8.4.3 Classification

To classify an event in a d-dimensional phase space as being either of signal or of background type, a local estimator of the probability that this event belongs to either class can be obtained from the foam's hyper-rectangular cells. The foams are created and filled based on samples of signal and background training events. For classification two possibilities are implemented. One foam can be used to separate the S/B probability density or two separate foams are created, one for the signal events and one for background events.

1) Separate signal and background foams

If the option `SigBgSeparate=True` is set, the method PDE-Foam treats the signal- and background distributions separately and performs the following steps to form the two foams and to calculate the classifier discriminator for a given event:

1. *Setup of training trees.* Two binary search trees are created and filled with the d-dimensional observable vectors of all signal and background training events, respectively.

2. *Initialisation phase.* Two independent foams for signal and background are created.

3. *Growing phase.* The growing is performed independently for the two foams. The density of events is estimated as the number of events found in the corresponding tree that are contained in the sampling box divided by the volume of the box (see `VolFrac` option). The geometries of the final foams reflect the distribution of the training samples: phase-space regions where the density of events is constant are combined in large cells, while in regions with large gradients in density many small cells are created.

4. *Filling phase.* Both for the signal and background foam each active cell is filled with the number of training events, n_S (signal) or n_B (background), contained in the corresponding cell volume, taking into account the event weights w_i: $n_S = \sum_{\text{sig. cell}} w_i$, $n_B = \sum_{\text{bg. cell}} w_i$.

5. *Evaluation phase.* The estimator for a given event is evaluated based on the number of events stored in the foam cells. The two corresponding foam cells that contain the event are found. The number of events (n_S and n_B) is read from the cells. The estimator $y_{\text{PDE-Foam}}(i)$ is then given by

$$y_{\text{PDE-Foam}}(i) = \frac{n_S/V_S}{\frac{n_B}{V_B} + \frac{n_S}{V_S}}\,, \tag{42}$$

where V_S and V_B are the respective cell volumes. The statistical error of the estimator is calculated as:

$$\sigma_{y_{\text{PDE-Foam}}}(i) = \sqrt{\left(\frac{n_S\sqrt{n_B}}{(n_S+n_B)^2}\right)^2 + \left(\frac{n_B\sqrt{n_S}}{(n_S+n_B)^2}\right)^2}\,. \tag{43}$$

Note that the so defined discriminant approximates the probability for an event from within the cell to be of class signal, if the weights are normalised such that the total number of weighted signal events is equal to the total number of weighted background events. This can be enforced with the normalisation mode "EqualNumEvents" (cf. Sec. 3.1.4).

Steps 1-4 correspond to the training phase of the method. Step 5 is performed during the testing phase. In contrast to the PDE-RS method the memory consumption and computation time for the testing phase does not depend on the number of training events.

Two separate foams for signal and background allow the foam algorithm to adapt the foam geometries to the individual shapes of the signal and background event distributions. It is therefore well suited for cases where the shapes of the two distributions are very different.

2) Single signal and background foam

If the option `SigBgSeparate=False` is set (default), the PDE-Foam method creates only one foam, which holds directly the estimator instead of the number of signal and background events. The differences with respect to separate signal and backgrounds foams are:

1. *Setup of training trees.* Fill only one binary search tree with both signal events and background events. This is possible, since the binary search tree has the information whether a event is of signal or background type.

2. *Initialisation phase.* Only one foam is created. The cells of this foam will contain the estimator $y_{\text{PDE-Foam}}(i)$ (see eq. (42)).

3. *Growing phase.* The splitting algorithm in this case minimises the variance of the estimator density $\sigma_\rho/\langle\rho\rangle$ across each cell. The estimator density ρ is sampled as the number of weighted signal events n_S over the total number of weighted events $n_S + n_B$ in a small box around the sampling points:
$$\rho = \frac{n_S}{n_S + n_B}\,\frac{1}{\texttt{VolFrac}} \tag{44}$$
In this case the geometries of the final foams reflect the distribution of the estimator density in the training sample: phase-space regions where the signal to background ratio is constant are combined in large cells, while in regions where the signal-to-background ratio changes rapidly many small cells are created.

4. *Filling phase.* Each active cell is filled with the estimator given as the ratio of weighted signal events to the total number of weighted events in the cell:
$$y_{\text{PDE-Foam}}(i) = \frac{n_S}{n_S + n_B}. \tag{45}$$
The statistical error of the estimator (43) also is stored in the cell.

5. *Evaluation phase.* The estimator for a given event is directly obtained as the discriminator that is stored in the cell which contains the event.

For the same total number of foam cells, the performance of the two implementations was found to be similar.

8.4.4 Regression

Two different methods are implemented for regression. In the first method, applicable for single targets only (*mono-target regression*), the target value is stored in each cell of the foam. In the second method, applicable to any number of targets (*multi-target regression*), the target values are stored in higher foam dimensions.

In **mono-target regression** the density used to form the foam is given by the mean target density in a given box. The procedure is as follows.

1. *Setup of training trees.* A binary search tree is filled with all training events.

2. *Growing phase.* One n_{var}-dimensional foam is formed: the density ρ is given by the mean target value $\langle t\rangle$ within the sampling box, divided by the box volume (given by the `VolFrac` option):
$$\rho = \frac{\langle t\rangle}{\texttt{VolFrac}} \equiv \frac{\sum_{i=1}^{N_{\text{box}}} t_i}{N_{\text{box}}\cdot\texttt{VolFrac}}\,, \tag{46}$$
where the sum is over all N_{box} events within the sampling box, and t_i is the target value of event i.

3. *Filling phase.* The cells are filled with their average target values, $\langle t \rangle = \sum_{i=1}^{N_{\text{box}}} t^{(i)}/N_{\text{box}}$.

4. *Evaluation phase.* Estimate the target value for a given test event: find the corresponding foam cell in which the test event is situated and read the average target value $\langle t \rangle$ from the cell.

For **multi-target regression** the target information is stored in additional foam dimensions. For a training sample with n_{var} (n_{tar}) input variables (regression targets), one would form a $(n_{\text{var}} + n_{\text{tar}})$-dimensional foam. To compute a target estimate for event i, one needs the coordinates of the cell centre $C(i, k)$ in each foam dimension k. The procedure is as follows.

1. *Setup of training trees.* A binary search tree is filled with all training events.

2. *Growing phase.* A $(n_{\text{var}} + n_{\text{tar}})$-dimensional foam is formed: the event density ρ is estimated by the number of events N_{box} within a predefined box of the dimension $(n_{\text{var}} + n_{\text{tar}})$, divided by the box volume, whereas the box volume is given by the `VolFrac` option

$$\rho = \frac{N_{\text{box}}}{\texttt{VolFrac}} \,. \tag{47}$$

3. *Filling phase.* Each foam cell is filled with the number of corresponding training events.

4. *Evaluation phase.* Estimate the target value for a given test event: find the N_{cells} foam cells in which the coordinates $(x_1, \ldots, x_{n_{\text{var}}})$ of the event vector are situated. Depending on the `TargetSelection` option, the target value t_k $(k = 1, \ldots, n_{\text{tar}})$ is

   (a) the coordinate of the cell centre in direction of the target dimension $n_{\text{var}} + k$ of the cell j $(j = 1, \ldots, N_{\text{cells}})$, which has the maximum event density

   $$t_k = C(j, n_{\text{var}} + k) \,, \tag{48}$$

   if `TargetSelection=Mpv`.

   (b) the mean cell centre in direction of the target dimension $n_{\text{var}} + k$ weighted by the event densities $d_{\text{ev}}(j)$ $(j = 1, \ldots, N_{\text{cells}})$ of the cells

   $$t_k = \frac{\sum_{j=1}^{N_{\text{cells}}} C(j, n_{\text{var}} + k) \cdot d_{\text{ev}}(j)}{\sum_{j=1}^{N_{\text{cells}}} d_{\text{ev}}(j)} \tag{49}$$

   if `TargetSelection=Mean`.

Kernel functions for regression

The kernel weighting methods have been implemented also for regression, taking into account the modified structure of the foam in case of multi-target regression.

8.4.5 Visualisation of the foam via projections to 2 dimensions

A projection function is used to visualise the foam in two dimensions. It is called via:

```
TH2D *proj = foam->Project2( dim1, dim2, "<options>", "<kernel>" );
```

Code Example 35: Call of the projection function. The first two arguments are the dimensions one wishes to project on, the third is a string identifier to specify quantity to plot (`nev`, `discr`, `rms`, `rms_ov_mean`, `MonoTargetRegression`, `MultiTargetRegression`), and the fourth argument chooses the kernel (`kNone`, `kGaus`).

For each active cell i the two-dimensional rectangular sub-space (dimensions `dim1` and `dim2`) is calculated and all bins contained in this sub-space are filled with the value $v(i)$ of cell i. This implies that in the case of more than two dimensions the values $v(i)$ in the dimensions that are not visible are summed. The filled value $v(i)$ depends on the given $\langle$`option`$\rangle$, which allows one to display all variables stored in the foam cells.

For the following description we define by $L(i,k)$ the length of the foam cell i in the dimension k of a d-dimensional foam. Symbolic, may $L(\text{Foam},k)$ be the scaled length of the entire foam in dimension k.

- `MultiTargetRegression`, `nev` – *projecting the number of events*
  These options apply to classification with separate signal and background foams and multi-target regression. The value $v(i)$ filled in the histogram is equal to the number of events $N_{\text{ev}}(i)$ stored in the foam cell i divided by the scaled two-dimensional cell area in dimension `dim1` and `dim2`
  $$v(i) = \frac{N_{\text{ev}}(i)}{L(i,\texttt{dim1}) \cdot L(i,\texttt{dim2}) \cdot L(\text{Foam},\texttt{dim1}) \cdot L(\text{Foam},\texttt{dim2})} \tag{50}$$

- `discr` – *projecting the discriminator*
  If the foam cells are filled with the discriminator, which is the case for classification with a single foam (`SigBgSeparate=False`), one can use this option. Here the value $v(i)$, filled in the histogram is equal to the discriminator $\text{Discr}(i)$ saved in cell i multiplied by the cell volume excluding the dimensions `dim1` and `dim2`
  $$v(i) = \text{Discr}(i) \prod_{\substack{k=1 \\ k\neq\texttt{dim1} \\ k\neq\texttt{dim2}}}^{d} L(i,k) \tag{51}$$
  This means that the average discriminator weighted with the cell size of the non-visualised dimensions is filled.

- `rms`, `rms_ov_mean` – *projection of cell variances*
  The variance (RMS) and the mean of the event distribution are saved in every foam cell.

If the option `rms` is used, the plotted cell value $v(i)$ is equal to the cell RMS. If the option `rms_ov_mean` is given, $v(i) = \text{RMS/Mean}$ is filled into the histogram.

- `MonoTargetRegression` – *projection of targets*
  If the foam cells are filled with targets by using the mono-target regression option in order to do regression (`MultiTargetRegression=False`), one can use this option. Here the value $v(i)$, filled in the histogram is equal to the target saved in cell i.

- `kNone`, `kGaus` – *Using kernels for projecting*
  Instead of filling the rectangle shaped cell areas into the histogram one can use the build-in kernel estimator to interpolate between the cell mean values in order to visualise the effect of the kernel to the foam. In this case the function performs a loop over all cells and fills the sum of the weighted cell values $v(i)$ into the histogram. See page 73 for more details on kernels included in PDE-Foam.

8.4.6 Performance

Like PDE-RS (see Sec. 8.3), this method is a powerful classification tool for problems with highly nonlinearly correlated observables. Furthermore PDE-Foam is a fast responding classifier, because of its limited number of cells, independent of the size of the training samples.

An exception is the multi-target regression with Gauss kernel because the time scales with the number of cells squared. Also the training can be slow, depending on the number of training events and number of cells one wishes to create.

8.5 k-Nearest Neighbour (k-NN) Classifier

Similar to PDE-RS (cf. Sec. 8.3), the k-nearest neighbour method compares an observed (test) event to reference events from a training data set [2]. However, unlike PDE-RS, which in its original form uses a fixed-sized multidimensional volume surrounding the test event, and in its augmented form resizes the volume as a function of the local data density, the k-NN algorithm is intrinsically adaptive. It searches for a fixed number of adjacent events, which then define a volume for the metric used. The k-NN classifier has best performance when the boundary that separates signal and background events has irregular features that cannot be easily approximated by parametric learning methods.

8.5.1 Booking options

The k-NN classifier is booked via the command:

Option	Array	Default	Predefined Values	Description
`nkNN`	–	`20`	–	Number of k-nearest neighbors
`BalanceDepth`	–	`6`	–	Binary tree balance depth
`ScaleFrac`	–	`0.8`	–	Fraction of events used to compute variable width
`SigmaFact`	–	`1`	–	Scale factor for sigma in Gaussian kernel
`Kernel`	–	`Gaus`	–	Use polynomial (=Poln) or Gaussian (=Gaus) kernel
`Trim`	–	`False`	–	Use equal number of signal and background events
`UseKernel`	–	`False`	–	Use polynomial kernel weight
`UseWeight`	–	`True`	–	Use weight to count kNN events
`UseLDA`	–	`False`	–	Use local linear discriminant - experimental feature

Option Table 14: Configuration options reference for MVA method: *k-NN*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured.

```
factory->BookMethod( Types::kKNN, "kNN", "<options>" );
```

Code Example 36: Booking of the k-NN classifier: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The configuration options for the k-NN classifier are listed in Option Table 14 (see also Sec. 6).

8.5.2 Description and implementation

The k-NN algorithm searches for k events that are closest to the test event. Closeness is thereby measured using a metric function. The simplest metric choice is the Euclidean distance

$$R = \left(\sum_{i=1}^{n_{\text{var}}} |x_i - y_i|^2 \right)^{\frac{1}{2}} . \tag{52}$$

where n_{var} is the number of input variables used for the classification, x_i are coordinates of an event from a training sample and y_i are variables of an observed test event. The k events with the smallest values of R are the *k-nearest neighbours*. The value of k determines the size of the neighbourhood for which a probability density function is evaluated. Large values of k do not capture the local behavior of the probability density function. On the other hand, small values of k cause statistical

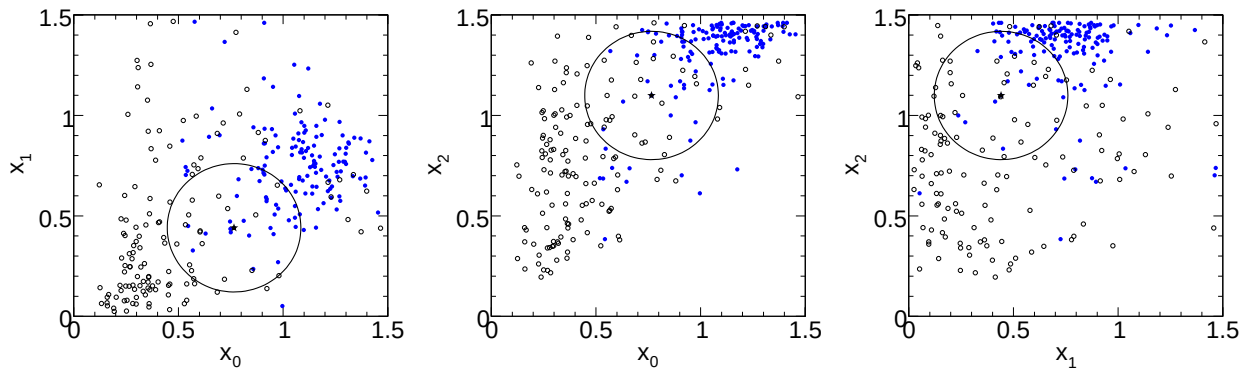

Figure 14: Example for the k-nearest neighbour algorithm in a three-dimensional space (i.e., for three discriminating input variables). The three plots are projections upon the two-dimensional coordinate planes. The full (open) circles are the signal (background) events. The k-NN algorithm searches for 20 nearest points in the *nearest neighborhood* (circle) of the query event, shown as a star. The nearest neighborhood counts 13 signal and 7 background points so that query event may be classified as a signal candidate.

fluctuations in the probability density estimate. A case study with real data suggests that values of k between 10 and 100 are appropriate and result in similar classification performance when the training sample contains hundreds of thousands of events (and $n_{\rm var}$ is of the order of a few variables).

The classification algorithm finds k-nearest training events around a query point

$$k = k_S + k_B \ , \tag{53}$$

where $k_{S(B)}$ is number of the signal (background) events in the training sample. The relative probability that the test event is of signal type is given by

$$P_S = \frac{k_S}{k_S + k_B} = \frac{k_S}{k} \ . \tag{54}$$

The choice of the metric governs the performance of the nearest neighbour algorithm. When input variables have different units a variable that has a wider distribution contributes with a greater weight to the Euclidean metric. This feature is compensated by rescaling the variables using a scaling fraction determined by the option `ScaleFrac`. Rescaling can be turned off by setting `ScaleFrac` to 0. The scaling factor applied to variable i is determined by the width w_i of the x_i distribution for the combined sample of signal and background events: w_i is the interval that contains the fraction `ScaleFrac` of x_i training values. The input variables are then rescaled by $1/w_i$, leading to the rescaled metric

$$R_{\rm rescaled} = \left(\sum_{i=1}^{d} \frac{1}{w_i^2} |x_i - y_i|^2 \right)^{\frac{1}{2}} \ . \tag{55}$$

Figure 14 shows an example of event classification with the k-nearest neighbour algorithm.[25]

[25] The number of training events shown has been greatly reduced to illustrate the principle of the algorithm. In a real application a typical k-NN training sample should be ample.

The output of the k-nearest neighbour algorithm can be interpreted as a probability that an event is of signal type, if the numbers (better: sum of event weights) of signal and background events in the training sample are equal. This can be enforced via the `Trim` option. If set training events of the overabundant type are randomly removed until parity is achieved.

Like (more or less) all TMVA classifiers, the k-nearest neighbour estimate suffers from statistical fluctuations in the training data. The typically high variance of the k-NN response is mitigated by adding a weight function that depends smoothly on the distance from a test event. The current k-NN implementation uses a polynomial kernel

$$W(x) = \begin{cases} (1-|x|^3)^3 & \text{if } |x| < 1 \,, \\ 0 & \text{otherwise} \,. \end{cases} \tag{56}$$

If R_k is the distance between the test event and the kth neighbour, the events are weighted according to the formula:

$$W_{S(B)} = \sum_{i=1}^{k_{S(B)}} W\left(\frac{R_i}{R_k}\right) \,, \tag{57}$$

where $k_{S(B)}$ is number of the signal (background) events in the neighbourhood. The weighted signal probability for the test event is then given by

$$P_S = \frac{W_S}{W_S + W_B} \,. \tag{58}$$

The kernel use is switched on/off by the option `UseKernel`.

Regression

The k-NN algorithm in TMVA also implements a simple multi-dimensional (multi-target) regression model. For a test event, the algorithm finds the k-nearest neighbours using the input variables, where each training event contains a regression value. The predicted regression value for the test event is the weighted average of the regression values of the k-nearest neighbours, cf. Eq. (39) on page 68.

8.5.3 Ranking

The present implementation of k-NN does not provide a ranking of the input variables.

8.5.4 Performance

The simplest implementation of the k-NN algorithm would store all training events in an array. The classification would then be performed by looping over all stored events and finding the k-nearest neighbours. As discussed in Sec. 4.2, such an implementation is impractical for large training samples. The k-NN algorithm therefore uses a *kd-tree* structure [12] that significantly improves the performance.

The TMVA implementation of the k-NN method is reasonably fast to allow classification of large data sets. In particular, it is faster than the adaptive PDE-RS method (cf. Sec. 8.3). Note that the k-NN method is not appropriate for problems where the number of input variables exceeds $n_{\text{var}} \gtrsim 10$. The neighbourood size R depends on n_{var} and the size of the training sample N as

$$R_N \propto \frac{1}{\sqrt[n_{\text{var}}]{N}} \,. \tag{59}$$

A large training set allows the algorithm to probe small-scale features that distinguish signal and background events.

8.6 H-Matrix discriminant

The origins of the H-Matrix approach dates back to works of Fisher and Mahalanobis in the context of Gaussian classifiers [22, 23]. It discriminates one class (signal) of a feature vector from another (background). The correlated elements of the vector are assumed to be Gaussian distributed, and the inverse of the covariance matrix is the *H-Matrix*. A multivariate χ^2 estimator is built that exploits differences in the mean values of the vector elements between the two classes for the purpose of discrimination.

The H-Matrix classifier as it is implemented in TMVA is equal or less performing than the Fisher discriminant (see Sec. 8.7), and has been only included for completeness.

8.6.1 Booking options

The H-Matrix discriminant is booked via the command:

```
factory->BookMethod( Types::kHMatrix, "HMatrix", "<options>" );
```

Code Example 37: Booking of the H-Matrix classifier: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

No specific options are defined for this method beyond those shared with all the other methods (cf. Option Table 9 on page 58).

8.6.2 Description and implementation

For an event i, each one χ^2 estimator ($\chi^2_{S(B)}$) is computed for signal (S) and background (B), using estimates for the sample means ($\overline{x}_{S(B),k}$) and covariance matrices ($C_{S(B)}$) obtained from the

training data

$$\chi^2_U(i) = \sum_{k,\ell=1}^{n_{\rm var}} (x_k(i) - \overline{x}_{U,k})\, C^{-1}_{U,k\ell}\, (x_\ell(i) - \overline{x}_{U,\ell}) \ , \tag{60}$$

where $U = S, B$. From this, the discriminant

$$y_H(i) = \frac{\chi^2_B(i) - \chi^2_S(i)}{\chi^2_B(i) + \chi^2_S(i)} \, , \tag{61}$$

is computed to discriminate between the signal and background classes.

8.6.3 Variable ranking

The present implementation of the H-Matrix discriminant does not provide a ranking of the input variables.

8.6.4 Performance

The TMVA implementation of the H-Matrix classifier has been shown to underperform in comparison with the corresponding Fisher discriminant (cf. Sec. 8.7), when using similar assumptions and complexity. It might therefore be considered to be depreciated.

8.7 Fisher discriminants (linear discriminant analysis)

In the method of Fisher discriminants [22] event selection is performed in a transformed variable space with zero linear correlations, by distinguishing the mean values of the signal and background distributions. The linear discriminant analysis determines an axis in the (correlated) hyperspace of the input variables such that, when projecting the output classes (signal and background) upon this axis, they are pushed as far as possible away from each other, while events of a same class are confined in a close vicinity. The linearity property of this classifier is reflected in the metric with which "far apart" and "close vicinity" are determined: the covariance matrix of the discriminating variable space.

8.7.1 Booking options

The Fisher discriminant is booked via the command:

```
factory->BookMethod( Types::kFisher, "Fisher", "<options>" );
```

Code Example 38: Booking of the Fisher discriminant: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

Option	Array	Default	Predefined Values	Description
`Method`	–	`Fisher`	`Fisher`, `Mahalanobis`	Discrimination method

Option Table 15: Configuration options reference for MVA method: *Fisher*. Values given are defaults. If predefined categories exist, the default category is marked by a '*star*'. The options in Option Table refopt:mva::methodbase on page pagerefopt:mva::methodbase can also be configured.

The configuration options for the Fisher discriminant are given in Option Table 15.

8.7.2 Description and implementation

The classification of the events in signal and background classes relies on the following characteristics: the overall sample means $\overline{x}_k$ for each input variable $k = 1, \ldots, n_{\rm var}$, the class-specific sample means $\overline{x}_{S(B),k}$, and total covariance matrix C of the sample. The covariance matrix can be decomposed into the sum of a *within-* (W) and a *between-class matrix* (B). They respectively describe the dispersion of events relative to the means of their own class (within-class matrix), and relative to the overall sample means (between-class matrix)[26].

The *Fisher coefficients*, F_k, are then given by

$$F_k = \frac{\sqrt{N_S N_B}}{N_S + N_B} \sum_{\ell=1}^{n_{\rm var}} W^{-1}_{k\ell} \left(\overline{x}_{S,\ell} - \overline{x}_{B,\ell}\right) , \tag{62}$$

where $N_{S(B)}$ are the number of signal (background) events in the training sample. The Fisher discriminant $y_{\rm Fi}(i)$ for event i is given by

$$y_{\rm Fi}(i) = F_0 + \sum_{k=1}^{n_{\rm var}} F_k x_k(i) \, . \tag{63}$$

The offset F_0 centers the sample mean $\overline{y}_{\rm Fi}$ of all $N_S + N_B$ events at zero.

Instead of using the within-class matrix, the Mahalanobis variant determines the Fisher coefficients

[26] The within-class matrix is given by

$$W_{k\ell} = \sum_{U=S,B} \langle x_{U,k} - \overline{x}_{U,k} \rangle \langle x_{U,\ell} - \overline{x}_{U,\ell} \rangle = C_{S,k\ell} + C_{B,k\ell} \, ,$$

where $C_{S(B)}$ is the covariance matrix of the signal (background) sample. The between-class matrix is obtained by

$$B_{k\ell} = \frac{1}{2} \sum_{U=S,B} \left(\overline{x}_{U,k} - \overline{x}_k\right) \left(\overline{x}_{U,\ell} - \overline{x}_\ell\right) ,$$

where $\overline{x}_{S(B),k}$ is the average of variable x_k for the signal (background) sample, and $\overline{x}_k$ denotes the average for the entire sample.

as follows [23]

$$F_k = \frac{\sqrt{N_S N_B}}{N_S + N_B} \sum_{\ell=1}^{n_{\text{var}}} C_{k\ell}^{-1} \left(\overline{x}_{S,\ell} - \overline{x}_{B,\ell} \right) , \tag{64}$$

where $C_{k\ell} = W_{k\ell} + B_{k\ell}$.

8.7.3 Variable ranking

The Fisher discriminant analysis aims at simultaneously maximising the between-class separation while minimising the within-class dispersion. A useful measure of the discrimination power of a variable is therefore given by the diagonal quantity B_{kk}/C_{kk}, which is used for the ranking of the input variables.

8.7.4 Performance

In spite of the simplicity of the classifier, Fisher discriminants can be competitive with likelihood and nonlinear discriminants in certain cases. In particular, Fisher discriminants are optimal for Gaussian distributed variables with linear correlations (cf. the standard toy example that comes with TMVA).

On the other hand, no discrimination at all is achieved when a variable has the same sample mean for signal and background, even if the shapes of the distributions are very different. Thus, Fisher discriminants often benefit from suitable transformations of the input variables. For example, if a variable $x \in [-1, 1]$ has a a signal distributions of the form x^2, and a uniform background distributions, their mean value is zero in both cases, leading to no separation. The simple transformation $x \to |x|$ renders this variable powerful for the use in a Fisher discriminant.

8.8 Linear discriminant analysis (LD)

The linear discriminant analysis provides data classification using a linear model, where *linear* refers to the discriminant function $y(\mathbf{x})$ being linear in the parameters β

$$y(\mathbf{x}) = \mathbf{x}^\top \beta + \beta_0 , \tag{65}$$

where β_0 (denoted the *bias*) is adjusted so that $y(\mathbf{x}) \geq 0$ for signal and $y(\mathbf{x}) < 0$ for background. It can be shown that this is equivalent to the Fisher discriminant, which seeks to maximise the ratio of between-class variance to within-class variance by projecting the data onto a linear subspace.

8.8.1 Booking options

The LD is booked via the command:

```
factory->BookMethod( Types::kLD, "LD" );
```

Code Example 39: Booking of the linear discriminant: the first argument is a predefined enumerator, the second argument is a user-defined string identifier. No method-specific options are available. See Sec. 3.1.5 for more information on the booking.

No specific options are defined for this method beyond those shared with all the other methods (cf. Option Table 9 on page 58).

8.8.2 Description and implementation

Assuming that there are $m+1$ parameters $\beta_0, \cdots, \beta_m$ to be estimated using a training set comprised of n events, the defining equation for β is

$$Y = X\beta \,, \tag{66}$$

where we have absorbed β_0 into the vector β and introduced the matrices

$$Y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \text{ and } X = \begin{pmatrix} 1 & x_{11} & \cdots & x_{1m} \\ 1 & x_{21} & \cdots & x_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n1} & \cdots & x_{nm} \end{pmatrix} \,, \tag{67}$$

where the constant column in X represents the bias β_0 and Y is composed of the target values with $y_i = 1$ if the ith event belongs to the signal class and $y_i = 0$ if the ith event belongs to the background class. Applying the method of least squares, we now obtain the *normal equations* for the classification problem, given by

$$X^T X\beta = X^T Y \Longleftrightarrow \beta = (X^T X)^{-1} X^T Y \,. \tag{68}$$

The transformation $(X^T X)^{-1} X^T$ is known as the *Moore-Penrose pseudo inverse* of X and can be regarded as a generalisation of the matrix inverse to non-square matrices. It requires that the matrix X has full rank.

If weighted events are used, this is simply taken into account by introducing a diagonal weight matrix W and modifying the normal equations as follows:

$$\beta = (X^T W X)^{-1} X^T W Y \,. \tag{69}$$

Considering two events $\mathbf{x}_1$ and $\mathbf{x}_2$ on the decision boundary, we have $y(\mathbf{x}_1) = y(\mathbf{x}_2) = 0$ and hence $(\mathbf{x}_1 - \mathbf{x}_2)^T \beta = 0$. Thus we see that the LD can be geometrically interpreted as determining the decision boundary by finding an orthogonal vector β.

8.8.3 Variable ranking

The present implementation of LD provides a ranking of the input variables based on the coefficients of the variables in the linear combination that forms the decision boundary. The order of importance of the discriminating variables is assumed to agree with the order of the absolute values of the coefficients.

8.8.4 Regression with LD

It is straightforward to apply the LD algorithm to linear regression by replacing the binary targets $y_i \in 0, 1$ in the training data with the measured values of the function which is to be estimated. The resulting function $y(\mathbf{x})$ is then the best estimate for the data obtained by least-squares regression.

8.8.5 Performance

The LD is optimal for Gaussian distributed variables with linear correlations (cf. the standard toy example that comes with TMVA) and can be competitive with likelihood and nonlinear discriminants in certain cases. No discrimination is achieved when a variable has the same sample mean for signal and background, but the LD can often benefit from suitable transformations of the input variables. For example, if a variable $x \in [-1, 1]$ has a signal distribution of the form x^2 and a uniform background distribution, their mean value is zero in both cases, leading to no separation. The simple transformation $x \rightarrow |x|$ renders this variable powerful for the use with LD.

8.9 Function discriminant analysis (FDA)

The common goal of all TMVA discriminators is to determine an optimal separating function in the multivariate space of all input variables. The Fisher discriminant solves this analytically for the linear case, while artificial neural networks, support vector machines or boosted decision trees provide nonlinear approximations with – in principle – arbitrary precision if enough training statistics is available and the chosen architecture is flexible enough.

The function discriminant analysis (FDA) provides an intermediate solution to the problem with the aim to solve relatively simple or partially nonlinear problems. The user provides the desired function with adjustable parameters via the configuration option string, and FDA fits the parameters to it, requiring the function value to be as close as possible to the real value (to 1 for signal and 0 for background in classification). Its advantage over the more involved and automatic nonlinear discriminators is the simplicity and transparency of the discrimination expression. A shortcoming is that FDA will underperform for involved problems with complicated, phase space dependent nonlinear correlations.

Option	Array	Default	Predefined Values	Description
`Formula`	–	`(0)`	–	The discrimination formula
`ParRanges`	–	`()`	–	Parameter ranges
`FitMethod`	–	`MINUIT`	`MC, GA, SA, MINUIT`	Optimisation Method
`Converger`	–	`None`	`None, MINUIT`	FitMethod uses Converger to improve result

Option Table 16: Configuration options reference for MVA method: *FDA*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. The input variables In the discriminator expression are denoted `x0`, `x1`, ...(until $n_{\text{var}} - 1$), where the number follows the order in which the variables have been registered with the Factory; coefficients to be determined by the fit must be denoted `(0)`, `(1)`, ...(the number of coefficients is free) in the formula; allowed is any functional expression that can be interpreted by a ROOT TFormula. See Code Example 41 for an example expression. The limits for the fit parameters (coefficients) defined in the formula expression are given with the syntax: "`(-1,3);(2,10);...`", where the first interval corresponds to parameter `(0)`. The converger allows to use (presently only) Minuit fitting in addition to Monte Carlo sampling or a Genetic Algorithm. More details on this combination are given in Sec. 6.5. The various fitters are configured using the options given in Tables 5, 6, 7 and 8, for MC, Minuit, GA and SA, respectively.

8.9.1 Booking options

FDA is booked via the command:

```
factory->BookMethod( Types::kFDA, "FDA", "<options>" );
```

Code Example 40: Booking of the FDA classifier: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The configuration options for the FDA classifier are listed in Option Table 16 (see also Sec. 6).

A typical option string could look as follows:

```
"Formula=(0)+(1)*x0+(2)*x1+(3)*x2+(4)*x3:\
 ParRanges=(-1,1);(-10,10);(-10,10);(-10,10);(-10,10):\
 FitMethod=MINUIT:\
 ErrorLevel=1:PrintLevel=-1:FitStrategy=2:UseImprove:UseMinos:SetBatch"
```

Code Example 41: FDA booking option example simulating a linear Fisher discriminant (cf. Sec. 8.7). The top line gives the discriminator expression, where the xi denote the input variables and the (j) denote the coefficients to be determined by the fit. Allowed are all standard functions and expressions, including the functions belonging to the ROOT TMath library. The second line determines the limits for the fit parameters, where the numbers of intervals given must correspond to the number of fit parameters defined. The third line defines the fitter to be used (here Minuit), and the last line is the fitter configuration.

8.9.2 Description and implementation

The parsing of the discriminator function employs ROOT's TFormula class, which requires that the expression complies with its rules (which are the same as those that apply for the `TTree::Draw` command). For simple formula with a single global fit solution, Minuit will be the most efficient fitter. However, if the problem is complicated, highly nonlinear, and/or has a non-unique solution space, more involved fitting algorithms may be required. In that case the Genetic Algorithm combined or not with a Minuit converger should lead to the best results. After fit convergence, FDA prints the fit results (parameters and estimator value) as well as the discriminator expression used on standard output. The smaller the estimator value, the better the solution found. The normalised estimator is given by

$$\begin{aligned} \textit{For classification:} \quad \mathcal{E} &= \frac{1}{W_S}\sum_{i=1}^{N_S}\left(F(\mathbf{x}_i)-1\right)^2 w_i + \frac{1}{W_B}\sum_{i=1}^{N_B} F^2(\mathbf{x}_i) w_i \ , \\ \textit{For regression:} \quad \mathcal{E} &= \frac{1}{W}\sum_{i=1}^{N}\left(F(\mathbf{x}_i)-\mathbf{t}_i\right)^2 w_i \ , \end{aligned} \tag{70}$$

where for classification the first (second) sum is over the signal (background) training events, and for regression it is over all training events, $F(\mathbf{x}_i)$ is the discriminator function, $\mathbf{x}_i$ is the tuple of the n_{var} input variables for event i, w_i is the event weight, $\mathbf{t}_i$ the tuple of training regression targets, $W_{S(B)}$ is the sum of all signal (background) weights in the training sample, and W the sum over all training weights.

8.9.3 Variable ranking

The present implementation of FDA does not provide a ranking of the input variables.

8.9.4 Performance

The FDA performance depends on the complexity and fidelity of the user-defined discriminator function. As a general rule, it should be able to reproduce the discrimination power of any linear

discriminant analysis. To reach into the nonlinear domain, it is useful to inspect the correlation profiles of the input variables, and add quadratic and higher polynomial terms between variables as necessary. Comparison with more involved nonlinear classifiers can be used as a guide.

8.10 Artificial Neural Networks (nonlinear discriminant analysis)

An Artificial Neural Network (ANN) is most generally speaking any simulated collection of interconnected neurons, with each neuron producing a certain response at a given set of input signals. By applying an external signal to some (input) neurons the network is put into a defined state that can be measured from the response of one or several (output) neurons. One can therefore view the neural network as a mapping from a space of input variables $x_1, \ldots, x_{\mathrm{n_{var}}}$ onto a one-dimensional (e.g. in case of a signal-versus-background discrimination problem) or multi-dimensional space of output variables $y_1, \ldots, y_{\mathrm{m_{var}}}$. The mapping is nonlinear if at least one neuron has a nonlinear response to its input.

In TMVA three neural network implementations are available to the user. The first was adapted from a FORTRAN code developed at the Université Blaise Pascal in Clermont-Ferrand,[27] the second is the ANN implementation that comes with ROOT. The third is a newly developed neural network (denoted *MLP*) that is faster and more flexible than the other two and is the recommended neural network to use with TMVA. All three neural networks are feed-forward multilayer perceptrons.

8.10.1 Booking options

The Clermont-Ferrand neural network

The Clermont-Ferrand neural network is booked via the command:

```
factory->BookMethod( Types::kCFMlpANN, "CF_ANN", "<options>" );
```

Code Example 42: Booking of the Clermont-Ferrand neural network: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The configuration options for the Clermont-Ferrand neural net are given in Option Table 17.

[27] The original Clermont-Ferrand neural network has been used for Higgs search analyses in ALEPH, and background fighting in rare B-decay searches by the BABAR Collaboration. For the use in TMVA the FORTRAN code has been converted to C++.

Option	Array	Default	Predefined Values	Description
`NCycles`	–	`3000`	–	Number of training cycles
`HiddenLayers`	–	`N,N-1`	–	Specification of hidden layer architecture

Option Table 17: Configuration options reference for MVA method: *CFMlpANN*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. See Sec. 8.10.3 for a description of the network architecture configuration.

Option	Array	Default	Predefined Values	Description
`NCycles`	–	`200`	–	Number of training cycles
`HiddenLayers`	–	`N,N-1`	–	Specification of hidden layer architecture (N stands for number of variables; any integers may also be used)
`ValidationFraction`	–	`0.5`	–	Fraction of events in training tree used for cross validation
`LearningMethod`	–	`Stochastic`	`Stochastic`, `Batch`, `SteepestDescent`, `RibierePolak`, `FletcherReeves`, `BFGS`	Learning method

Option Table 18: Configuration options reference for MVA method: *TMlpANN*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. See Sec. 8.10.3 for a description of the network architecture configuration.

The ROOT neural network (class TMultiLayerPerceptron)

This neural network interfaces the ROOT class `TMultiLayerPerceptron` and is booked through the Factory via the command line:

```
factory->BookMethod( Types::kTMlpANN, "TMlp_ANN", "<options>" );
```

Code Example 43: Booking of the ROOT neural network: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. See Sec. 3.1.5 for more information on the booking.

The configuration options for the ROOT neural net are given in Option Table 18.

The MLP neural network

The MLP neural network is booked through the Factory via the command line:

```
factory->BookMethod( Types::kMLP, "MLP_ANN", "<options>" );
```

Code Example 44: Booking of the MLP neural network: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the options string. See Sec. 3.1.5 for more information on the booking.

The configuration options for the MLP neural net are given in Option Table 19. The TMVA implementation of MLP supports random and importance event sampling. With event sampling enabled, only a fraction (set by the option `Sampling`) of the training events is used for the training of the MLP. Values in the interval $[0, 1]$ are possible. If the option `SamplingImportance` is set to 1, the events are selected randomly, while for a value below 1 the probability for the same events to be sampled again depends on the training performance achieved for classification or regression. If for a given set of events the training leads to a decrease of the error of the test sample, the probability for the events of being selected again is multiplied with the factor given in `SamplingImportance` and thus decreases. In the case of an increased error of the test sample, the probability for the events to be selected again is divided by the factor `SamplingImportance` and thus increases. The probability for an event to be selected is constrained to the interval $[0, 1]$. For each set of events, the importance sampling described above is performed together with the overtraining test.

Event sampling is performed until the fraction specified by the option `SamplingEpoch` of the total number of epochs (`NCycles`) has been reached. Afterwards, all available training events are used for the training. Event sampling can be turned on and off for training and testing events individually with the options `SamplingTraining` and `SamplingTesting`.

The aim of random and importance sampling is foremost to speed-up the training for large training samples. As a side effect, random or importance sampling may also increase the robustness of the training algorithm with respect to convergence in a local minimum.

Since it is typically not known beforehand how many epochs are necessary to achieve a sufficiently good training of the neural network, a convergence test can be activated by setting `ConvergenceTests` to a value above 0. This value denotes the number of subsequent convergence tests which have to fail (i.e. no improvement of the estimator larger than `ConvergenceImprove`) to consider the training to be complete. Convergence tests are performed at the same time as overtraining tests. The test frequency is given by the parameter `TestRate`.

It is recommended to set the option `VarTransform=Norm`, such that the input (and in case of regression the output as well) is normalised to the interval $[-1, 1]$.

Option	Array	Default	Predefined Values	Description
NCycles	–	500	–	Number of training cycles
HiddenLayers	–	N,N-1	–	Specification of hidden layer architecture
NeuronType	–	sigmoid	linear, sigmoid, tanh, radial	Neuron activation function type
NeuronInputType	–	sum	sum, sqsum, abssum	Neuron input function type
TrainingMethod	–	BP	BP, GA, BFGS	Train with Back-Propagation (BP), BFGS Algorithm (BFGS), or Genetic Algorithm (GA - slower and worse)
LearningRate	–	0.02	–	ANN learning rate parameter
DecayRate	–	0.01	–	Decay rate for learning parameter
TestRate	–	10	–	Test for overtraining performed at each #th epochs
Sampling	–	1	–	Only 'Sampling' (randomly selected) events are trained each epoch
SamplingEpoch	–	1	–	Sampling is used for the first 'SamplingEpoch' epochs, afterwards, all events are taken for training
SamplingImportance	–	1	–	The sampling weights of events in epochs which successful (worse estimator than before) are multiplied with SamplingImportance, else they are divided.
SamplingTraining	–	True	–	The training sample is sampled
SamplingTesting	–	False	–	The testing sample is sampled
ResetStep	–	50	–	How often BFGS should reset history
Tau	–	3	–	LineSearch size step
BPMode	–	sequential	sequential, batch	Back-propagation learning mode: sequential or batch
BatchSize	–	-1	–	Batch size: number of events/batch, only set if in Batch Mode, -1 for BatchSize=number_of_events
ConvergenceImprove	–	0	–	Minimum improvement which counts as improvement (<0 means automatic convergence check is turned off)
ConvergenceTests	–	-1	–	Number of steps (without improvement) required for convergence (<0 means automatic convergence check is turned off)

Option Table 19: Configuration options reference for MVA method: *MLP*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. See Sec. 8.10.3 for a description of the network architecture configuration.

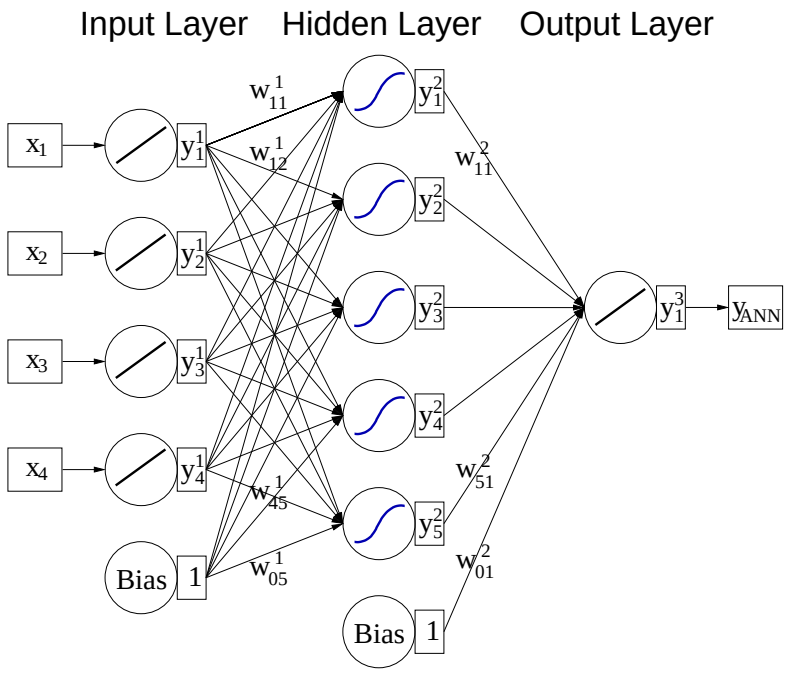

Figure 15: Multilayer perceptron with one hidden layer.

8.10.2 Description and implementation

The behaviour of an artificial neural network is determined by the layout of the neurons, the weights of the inter-neuron connections, and by the response of the neurons to the input, described by the *neuron response function* ρ.

Multilayer Perceptron

While in principle a neural network with n neurons can have n^2 directional connections, the complexity can be reduced by organising the neurons in layers and only allowing direct connections from a given layer to the following layer (see Fig. 15). This kind of neural network is termed *multi-layer perceptron*; all neural net implementations in TMVA are of this type. The first layer of a multilayer perceptron is the input layer, the last one the output layer, and all others are *hidden* layers. For a classification problem with n_{var} input variables the input layer consists of n_{var} neurons that hold the input values, $x_1, \ldots, x_{n_{\text{var}}}$, and one neuron in the output layer that holds the output variable, the neural net estimator y_{ANN}.

For a regression problem the network structure is similar, except that for multi-target regression each of the targets is represented by one output neuron. A weight is associated to each directional connection between the output of one neuron and the input of another neuron. When calculating the input value to the response function of a neuron, the output values of all neurons connected to the given neuron are multiplied with theses weights.

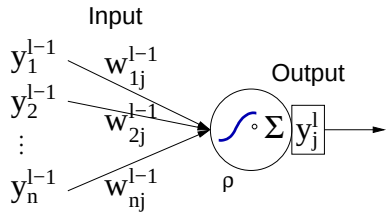

Figure 16: Single neuron j in layer ℓ with n input connections. The incoming connections carry a weight of $w_{ij}^{(l-1)}$.

Neuron response function

The neuron response function ρ maps the neuron input $i_1, \ldots, i_n$ onto the neuron output (Fig. 16). Often it can be separated into a $\mathcal{R}^n \mapsto \mathcal{R}$ *synapse function* κ, and a $\mathcal{R} \mapsto \mathcal{R}$ *neuron activation function* α, so that $\rho = \alpha \circ \kappa$. The functions κ and α can have the following forms:

$$\kappa: \ (y_1^{(\ell)}, .., y_n^{(\ell)} | w_{0j}^{(\ell)}, .., w_{nj}^{(\ell)}) \rightarrow \begin{cases} w_{0j}^{(\ell)} + \sum_{i=1}^{n} y_i^{(\ell)} w_{ij}^{(\ell)} & \textit{Sum}, \\ w_{0j}^{(\ell)} + \sum_{i=1}^{n} \left(y_i^{(\ell)} w_{ij}^{(\ell)} \right)^2 & \textit{Sum of squares}, \\ w_{0j}^{(\ell)} + \sum_{i=1}^{n} |y_i^{(\ell)} w_{ij}^{(\ell)}| & \textit{Sum of absolutes}, \end{cases} \tag{71}$$

$$\alpha: \ x \rightarrow \begin{cases} x & \textit{Linear}, \\ \dfrac{1}{1+e^{-kx}} & \textit{Sigmoid}, \\ \dfrac{e^x - e^{-x}}{e^x + e^{-x}} & \textit{Tanh}, \\ e^{-x^2/2} & \textit{Radial}. \end{cases} \tag{72}$$

8.10.3 Network architecture

The number of hidden layers in a network and the number of neurons in these layers are configurable via the option `HiddenLayers`. For example the configuration `"HiddenLayers=N-1,N+10,3"` creates a network with three hidden layers, the first hidden layer with $n_{\text{var}} - 1$ neurons, the second with $n_{\text{var}} + 10$ neurons, and the third with 3 neurons.

When building a network two rules should be kept in mind. The first is the theorem by Weierstrass, which if applied to neural nets, ascertains that for a multilayer perceptron a single hidden layer is sufficient to approximate a given continuous correlation function to any precision, provided that a sufficiently large number of neurons is used in the hidden layer. If the available computing power and the size of the training data sample suffice, one can increase the number of neurons in the hidden layer until the optimal performance is reached.

It is likely that the same performance can be achieved with a network of more than one hidden layer and a potentially much smaller total number of hidden neurons. This would lead to a shorter training time and a more robust network.

8.10.4 Training of the neural network

Back-propagation (BP)

The most common algorithm for adjusting the weights that optimise the classification performance of a neural network is the so-called *back propagation.* It belongs to the family of supervised learning methods, where the desired output for every input event is known. Back propagation is used by all neural networks in TMVA. The output of a network (here for simplicity assumed to have a single hidden layer with a Tanh activation function, and a linear activation function in the output layer) is given by

$$y_{\mathrm{ANN}} = \sum_{j=1}^{n_{\mathrm{h}}} y_j^{(2)} w_{j1}^{(2)} = \sum_{j=1}^{n_{\mathrm{h}}} \tanh\left(\sum_{i=1}^{n_{\mathrm{var}}} x_i w_{ij}^{(1)}\right) \cdot w_{j1}^{(2)} \,, \tag{73}$$

where n_{var} and n_{h} are the number of neurons in the input layer and in the hidden layer, respectively, $w_{ij}^{(1)}$ is the weight between input-layer neuron i and hidden-layer neuron j, and $w_{j1}^{(2)}$ is the weight between the hidden-layer neuron j and the output neuron. A simple sum was used in Eq. (73) for the synapse function κ.

During the learning process the network is supplied with N training events $\mathbf{x}_a = (x_1, \ldots, x_{n_{\mathrm{var}}})_a$, $a = 1, \ldots, N$. For each training event a the neural network output $y_{\mathrm{ANN},a}$ is computed and compared to the desired output $\hat{y}_a \in \{1, 0\}$ (in classification 1 for signal events and 0 for background events). An *error function* E, measuring the agreement of the network response with the desired one, is defined by

$$E(\mathbf{x}_1, \ldots, \mathbf{x}_N|\mathbf{w}) = \sum_{a=1}^{N} E_a(\mathbf{x}_a|\mathbf{w}) = \sum_{a=1}^{N} \frac{1}{2} \left(y_{\mathrm{ANN},a} - \hat{y}_a\right)^2 \,, \tag{74}$$

where $\mathbf{w}$ denotes the ensemble of adjustable weights in the network. The set of weights that minimises the error function can be found using the method of *steepest* or *gradient descent*, provided that the neuron response function is differentiable with respect to the input weights. Starting from a random set of weights $\mathbf{w}^{(\rho)}$ the weights are updated by moving a small distance in $\mathbf{w}$-space into the direction $-\boldsymbol{\nabla}_{\mathbf{w}} E$ where E decreases most rapidly

$$\mathbf{w}^{(\rho+1)} = \mathbf{w}^{(\rho)} - \eta \boldsymbol{\nabla}_{\mathbf{w}} E \,, \tag{75}$$

where the positive number η is the *learning rate.*

The weights connected with the output layer are updated by

$$\Delta w_{j1}^{(2)} = -\eta \sum_{a=1}^{N} \frac{\partial E_a}{\partial w_{j1}^{(2)}} = -\eta \sum_{a=1}^{N} \left(y_{\mathrm{ANN},a} - \hat{y}_a\right) y_{j,a}^{(2)} \,, \tag{76}$$

and the weights connected with the hidden layers are updated by

$$\Delta w_{ij}^{(1)} = -\eta \sum_{a=1}^{N} \frac{\partial E_a}{\partial w_{ij}^{(1)}} = -\eta \sum_{a=1}^{N} (y_{\mathrm{ANN},a} - \hat{y}_a)\, y_{j,a}^{(2)} (1 - y_{j,a}^{(2)}) w_{j1}^{(2)} x_{i,a}\,, \tag{77}$$

where we have used $\tanh' x = \tanh x(1 - \tanh x)$. This method of training the network is denoted *bulk learning*, since the sum of errors of all training events is used to update the weights. An alternative choice is the so-called *online learning*, where the update of the weights occurs at each event. The weight updates are obtained from Eqs. (76) and (77) by removing the event summations. In this case it is important to use a well randomised training sample. Online learning is the learning method implemented in TMVA.

BFGS

The Broyden-Fletcher-Goldfarb-Shannon (BFGS) method [24] differs from back propagation by the use of second derivatives of the error function to adapt the synapse weight by an algorithm which is composed of four main steps.

1. Two vectors, D and Y are calculated. The vector of weight changes D represents the evolution between one iteration of the algorithm $(k-1)$ to the next (k). Each synapse weight corresponds to one element of the vector. The vector Y is the vector of gradient errors.

$$D_i^{(k)} = w_i^{(k)} - w_i^{(k-1)}\,, \tag{78}$$
$$Y_i^{(k)} = g_i^{(k)} - g_i^{(k-1)}\,, \tag{79}$$

   where i is the synapse index, g_i is the i-th synapse gradient,[28] w_i is the weight of the i-th synapse, and k denotes the iteration counter.

2. Approximate the inverse of the Hessian matrix, H^{-1}, at iteration k by

$$H^{-1(k)} = \frac{D \cdot D^T \cdot (1 + Y^T \cdot H^{-1(k-1)} \cdot Y)}{Y^T \cdot D} - D \cdot Y^T \cdot H + H \cdot Y \cdot D^T + H^{-1(k-1)}\,, \tag{80}$$

   where superscripts (k) are implicit for D and Y.

3. Estimate the vector of weight changes by

$$D^{(k)} = -H^{-1(k)} \cdot Y^{(k)}\,. \tag{81}$$

4. Compute a new vector of weights by applying a *line search* algorithm. In the line search the error function is locally approximated by a parabola. The algorithm evaluates the second derivatives and determines the point where the minimum of the parabola is expected. The total error is evaluated for this point. The algorithm then evaluates points along the line

[28] The synapse gradient is estimated in the same way as in the BP method (with initial gradient and weights set to zero).

defined by the direction of the gradient in weights space to find the absolute minimum. The weights at the minimum are used for the next iteration. The learning rate can be set With the option `Tau`. The learning parameter, which defines by how much the weights are changed in one epoch along the line where the minimum is suspected, is multiplied with the learning rate as long as the training error of the neural net with the changed weights is below the one with unchanged weights. If the training error of the changed neural net were already larger for the initial learning parameter, it is divided by the learning rate until the training error becomes smaller. The iterative and approximate calculation of $H^{-1(k)}$ turns less accurate with an increasing number of iterations. The matrix is therefore reset to the unit matrix every `ResetStep` steps.

The advantage of the BFGS method compared to BG is the smaller number of iterations. However, because the computing time for one iteration is proportional to the squared number of synapses, large networks are particularly penalised.

8.10.5 Variable ranking

The MLP neural network implements a variable ranking that uses the sum of the weights-squared of the connections between the variable's neuron in the input layer and the first hidden layer. The importance I_i of the input variable i is given by

$$I_i = \overline{x}_i^2 \sum_{j=1}^{n_\mathrm{h}} \left(w_{ij}^{(1)} \right)^2 , \qquad i = 1, \ldots, n_\mathrm{var} , \tag{82}$$

where $\overline{x}_i$ is the sample mean of input variable i.

8.10.6 Performance

In the tests we have carried out so far, the MLP and ROOT networks performed equally well, however with a clear speed advantage for the MLP. The Clermont-Ferrand neural net exhibited worse classification performance in these tests, which is partly due to the slow convergence of its training (at least 10k training cycles are required to achieve approximately competitive results).

8.11 Support Vector Machine (SVM)

In the early 1960s a linear support vector method has been developed for the construction of separating hyperplanes for pattern recognition problems [37, 38]. It took 30 years before the method was generalised to nonlinear separating functions [39, 40] and for estimating real-valued functions (regression) [41]. At that moment it became a general purpose algorithm, performing classification and regression tasks which can compete with neural networks and probability density estimators. Typical applications of SVMs include text categorisation, character recognition, bio-informatics and face detection.

Option	Array	Default	Predefined Values	Description
`C`	–	`1`	–	Cost parameter
`Tol`	–	`0.01`	–	Tolerance parameter
`MaxIter`	–	`1000`	–	Maximum number of training loops
`NSubSets`	–	`1`	–	Number of training subsets
`Gamma`	–	`1`	–	RBF kernel parameter: Gamma

Option Table 20: Configuration options reference for MVA method: *SVM*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. Definition of the kernel function: `Linear` is $K(\vec{x},\vec{y}) = \vec{x}\cdot\vec{y}$ (no extra parameters), `Polynomial` is $K(\vec{x},\vec{y}) = (\vec{x}\cdot\vec{y}+\theta)^d$, `Gauss` is $K(\vec{x},\vec{y}) = \exp\left(-|\vec{x}-\vec{y}|^2/2\sigma^2\right)$, and `Sigmoid` corresponds to $K(\vec{x},\vec{y}) = \tanh\left(\kappa(\vec{x}\cdot\vec{y})+\theta\right)$

The main idea of the SVM approach to classification problems is to build a hyperplane that separates signal and background *vectors* (events) using only a minimal subset of all training vectors (*support vectors*). The position of the hyperplane is obtained by maximizing the margin (distance) between it and the support vectors. The extension to nonlinear SVMs is performed by mapping the input vectors onto a higher dimensional feature space in which signal and background events can be separated by a linear procedure using an optimally separating hyperplane. The use of kernel functions eliminates thereby the explicit transformation to the feature space and simplifies the computation.

The implementation of the newly introduced regression is similar to the approach in classification. It also maps input data into higher dimensional space using previously chosen support vectors. Instead of separating events of two types, it determines the hyperplane with events of the same value (which is equal to the *mean* from all training events). The final value is estimated based on the distance to the hyperplane which is computed by the selected kernel function.

8.11.1 Booking options

The SVM classifier is booked via the command:

```
factory->BookMethod( TMVA::Types::kSVM, "SVM", "<options>" );
```

Code Example 45: Booking of the SVM classifier: the first argument is a unique type enumerator, the second is a user-defined name which must be unique among all booked classifiers, and the third argument is the configuration option string. Individual options are separated by a ':'. For options that are not set in the string default values are used. See Sec. 3.1.5 for more information on the booking.

The configuration options for the SVM classifier are given in Option Table 20.

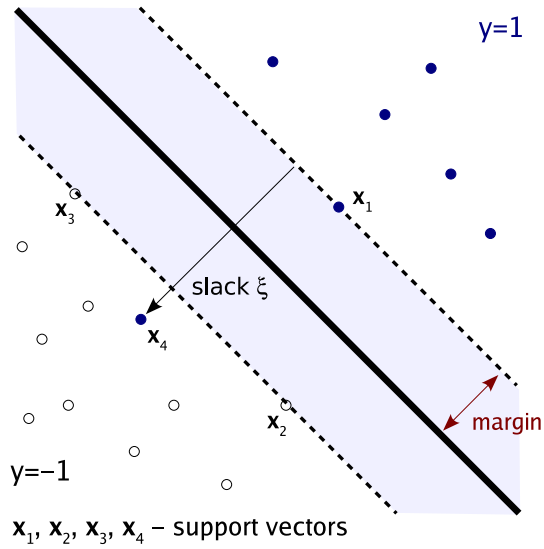

Figure 17: Hyperplane classifier in two dimensions. The vectors (events) $\mathbf{x}_{1-4}$ define the hyperplane and margin, i.e., they are the support vectors.

8.11.2 Description and implementation

A detailed description of the SVM formalism can be found, for example, in Ref. [42]. Here only a brief introduction of the TMVA implementation is given.

Linear SVM

Consider a simple two-class classifier with oriented hyperplanes. If the training data is linearly separable, a vector-scalar pair $(\vec{w}, b)$ can be found that satisfies the constraints

$$y_i(\vec{x}_i \cdot \vec{w} + b) - 1 \geq 0\,, \qquad \forall_i\,, \tag{83}$$

where $\vec{x}_i$ are the input vectors, y_i the desired outputs ($y_i = \pm 1$), and where the pair $(\vec{w}, b)$ defines a hyperplane. The decision function of the classifier is $f(\vec{x}_i) = \mathrm{sign}(\vec{x}_i \cdot \vec{w} + b)$, which is $+1$ for all points on one side of the hyperplane and -1 for the points on the other side.

Intuitively, the classifier with the largest margin will give better separation. The margin for this linear classifier is just $2/|\vec{w}|$. Hence to maximise the margin, one needs to minimise the *cost function* $W = |\vec{w}|^2/w$ with the constraints from Eq. (83).

At this point it is beneficial to consider the significance of different input vectors $\vec{x}_i$. The training events lying on the margins, which are called the support vectors (SV), are the events that contribute to defining the decision boundary (see Fig. 17). Hence if the other events are removed from the training sample and the classifier is retrained on the remaining events, the training will result in

the same decision boundary. To solve the constrained quadratic optimisation problem, we first reformulate it in terms of a Lagrangian

$$\mathcal{L}(\vec{w}, b, \vec{\alpha}) = \frac{1}{2}\,|\vec{w}|^2 - \sum_i \alpha_i \left(y_i \left((\vec{x}_i \cdot \vec{w}) + b\right) - 1\right) \tag{84}$$

where $\alpha_i \geq 0$ and the condition from Eq. (83) must be fulfilled. The Lagrangian $\mathcal{L}$ is minimised with respect to $\vec{w}$ and b and maximised with respect to $\vec{\alpha}$. The solution has an expansion in terms of a subset of input vectors for which $\alpha_i \neq 0$ (the support vectors):

$$\vec{w} = \sum_i \alpha_i y_i \vec{x}_i \,, \tag{85}$$

because $\partial\mathcal{L}/\partial b = 0$ and $\partial\mathcal{L}/\partial\vec{w} = 0$ hold at the extremum. The optimisation problem translates to finding the vector $\vec{\alpha}$ which maximises

$$\mathcal{L}(\vec{\alpha}) = \sum_i \alpha_i - \frac{1}{2}\sum_{ij} \alpha_i \alpha_j y_i y_j \vec{x}_i \cdot \vec{x}_j \,. \tag{86}$$

Both the optimisation problem and the final decision function depend only on scalar products between input vectors, which is a crucial property for the generalisation to the nonlinear case.

Nonseparable data

The above algorithm can be extended to non-separable data. The classification constraints in Eq. (83) are modified by adding a "slack" variable ξ_i to it ($\xi_i = 0$ if the vector is properly classified, otherwise ξ_i is the distance to the decision hyperplane)

$$y_i(\vec{x}_i \cdot \vec{w} + b) - 1 + \xi_i \geq 0, \qquad \xi_i \geq 0\,, \quad \forall_i\,. \tag{87}$$

This admits a certain amount of misclassification. The training algorithm thus minimises the modified cost function

$$W = \frac{1}{2}\,|\vec{w}|^2 + C\sum_i \xi_i \,, \tag{88}$$

describing a trade-off between margin and misclassification. The cost parameter `C` sets the scale by how much misclassification increases the cost function (see Tab. 20).

Nonlinear SVM

The SVM formulation given above can be further extended to build a nonlinear SVM which can classify nonlinearly separable data. Consider a function $\Phi:\ \mathrm{R}^{n_{\mathrm{var}}} \to \mathcal{H}$, which maps the training data from $\mathrm{R}^{n_{\mathrm{var}}}$, where n_{var} is the number of discriminating input variables, to some higher dimensional space $\mathcal{H}$. In the $\mathcal{H}$ space the signal and background events can be linearly separated so that the linear SVM formulation can be applied. We have seen in Eq. (86) that event variables only

appear in the form of scalar products $\vec{x}_i \cdot \vec{x}_j$, which become $\Phi(\vec{x}_i) \cdot \Phi(\vec{x}_j)$ in the higher dimensional feature space $\mathcal{H}$. The latter scalar product can be approximated by a kernel function

$$K(\vec{x}_i, \vec{x}_j) \approx \Phi(\vec{x}_i) \cdot \Phi(\vec{x}_j)\,, \tag{89}$$

which avoids the explicit computation of the mapping function $\Phi(\vec{x})$. This is desirable because the exact form of $\Phi(\vec{x})$ is hard to derive from the training data. Most frequently used kernel functions are

$$\begin{aligned} K(\vec{x}, \vec{y}) &= (\vec{x} \cdot \vec{y} + \theta)^d && \textit{Polynomial}, \\ K(\vec{x}, \vec{y}) &= \exp\left(-\left|\vec{x} - \vec{y}\right|^2 / 2\sigma^2\right) && \textit{Gaussian}, \\ K(\vec{x}, \vec{y}) &= \tanh\left(\kappa(\vec{x} \cdot \vec{y}) + \theta\right) && \textit{Sigmoidal}. \end{aligned} \tag{90}$$

It was shown in Ref. [41] that a suitable function kernel must fulfill Mercer's condition

$$\int K(\vec{x}, \vec{y}) g(\vec{x}) g(\vec{y}) d\vec{x} d\vec{y} \geq 0\,, \tag{91}$$

for any function g such that $\int g^2(\vec{x}) d\vec{x}$ is finite. While Gaussian and polynomial kernels are known to comply with Mercer's condition, this is not strictly the case for sigmoidal kernels. To extend the linear methodology to nonlinear problems one substitutes $\vec{x}_i \cdot \vec{x}_j$ by $K(\vec{x}_i, \vec{x}_j)$ in Eq. (86). Due to Mercer's conditions on the kernel, the corresponding optimisation problem is a well defined convex quadratic programming problem with a global minimum. This is an advantage of SVMs compared to neural networks where local minima occur.

For regression problems, the same algorithm is used as for classification with the exception that instead of dividing events based on their type (signal/background), it separates them based on the value (larger/smaller than average). In the end, it does not return the sigmoid of the distance between the event and the hyperplane, but the distance itself – increased by the average target value.

Implementation

The TMVA implementation of the Support Vector Machine follows closely the description given in the literature. It employs a sequential minimal optimisation (SMO) [43] to solve the quadratic problem. Acceleration of the minimisation is achieved by dividing a set of vectors into smaller subsets [44]. The number of training subsets is controlled by option `NSubSets`. The SMO method drives the subset selection to the extreme by selecting subsets of two vectors (for details see Ref. [42]). The pairs of vectors are chosen, using heuristic rules, to achieve the largest possible improvement (minimisation) per step. Because the working set is of size two, it is straightforward to write down the analytical solution. The minimisation procedure is repeated recursively until the minimum is found. The SMO algorithm has proven to be significantly faster than other methods and has become the most common minimisation method used in SVM implementations. The precision of the minimisation is controlled by the tolerance parameter `Tol` (see Tab. 20). The SVM training time can be reduced by increasing the tolerance. Most classification problems should be solved with less then 1000 training iterations. Interrupting the SVM algorithm using the option `MaxIter` may thus be helpful when optimising the SVM training parameters. `MaxIter` can be released for the final classifier training.

8.11.3 Variable ranking

The present implementation of the SVM classifier does not provide a ranking of the input variables.

8.11.4 Performance

The TMVA SVM algorithm comes with linear, polynomial, Gaussian and sigmoidal kernel functions. With sufficient training statistics, the Gaussian kernel allows to approximate any separating function in the input space. It is crucial for the performance of the SVM to appropriately tune the kernel parameters and the cost parameter `C`. In case of a Gaussian, the kernel is tuned via option `Gamma` which is related to the width σ by $\Gamma = 1/(2\sigma^2)$. The optimal tuning of these parameters is specific to the problem and must be done by the user.

The SVM training time scales with n^2, where n is the number of vectors (events) in the training data set. The user is therefore advised to restrict the sample size during the first rough scan of the kernel parameters. Also increasing the minimisation tolerance helps to speed up the training.

SVM is a nonlinear general purpose classification and regression algorithm with a performance similar to neural networks (Sec. 8.10) or to a multidimensional likelihood estimator (Sec. 8.3).

8.12 Boosted Decision and Regression Trees

A *decision (regression) tree* (BDT)[29] is a binary tree structured classifier (regressor) similar to the one sketched in Fig. 18. Repeated left/right (yes/no) decisions are taken on one single variable at a time until a stop criterion is fulfilled. The phase space is split this way into many regions that are eventually classified as signal or background, depending on the majority of training events that end up in the final *leaf* node. In case of *regression trees*, each output node represents a specific value of the target variable.[30] The boosting (see Sec. 7) of a decision (regression) tree extends this concept from one tree to several trees which form a *forest*. The trees are derived from the same training ensemble by reweighting events, and are finally combined into a single classifier (regressor) which is given by a (weighted) average of the individual decision (regression) trees. Boosting stabilizes the response of the decision trees with respect to fluctuations in the training sample and is able to considerably enhance the performance w.r.t. a single tree. In the following, we will use the term *decision tree* for both, *decision-* and *regression trees* and we refer to regression trees only if both types are treated differently.

8.12.1 Booking options

The BDT classifier is booked via the command:

[29] We use the acronym BDT for *decision* as well as *regression* trees.
[30] The target variable is the variable the regression "function" is trying to estimate.

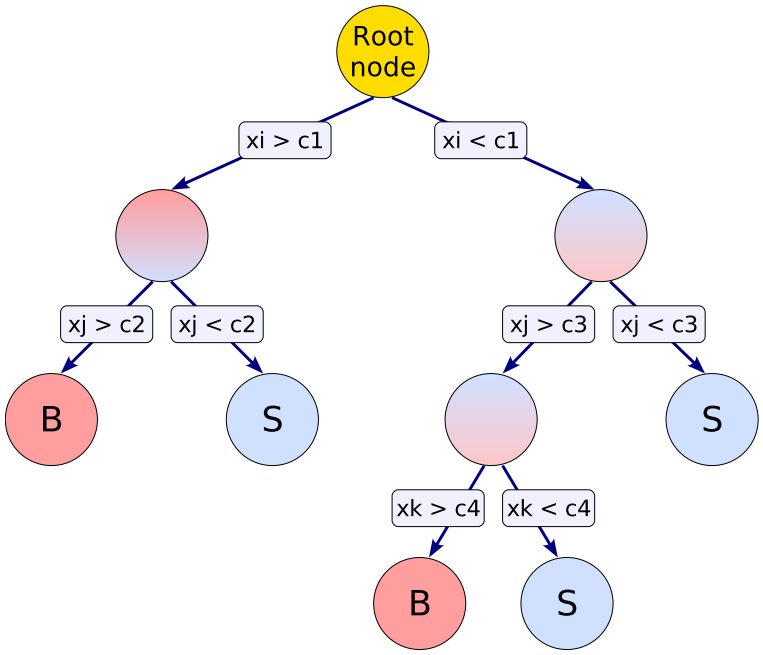

Figure 18: Schematic view of a decision tree. Starting from the root node, a sequence of binary splits using the discriminating variables x_i is applied to the data. Each split uses the variable that at this node gives the best separation between signal and background when being cut on. The same variable may thus be used at several nodes, while others might not be used at all. The leaf nodes at the bottom end of the tree are labeled "S" for signal and "B" for background depending on the majority of events that end up in the respective nodes. For regression trees, the node splitting is performed on the variable that gives the maximum decrease in the average squared error when attributing a constant value of the target variable as output of the node, given by the average of the training events in the corresponding (leaf) node (see Sec. 8.12.3).

```
factory->BookMethod( Types::kBDT, "BDT", "<options>" );
```

Code Example 46: Booking of the BDT classifier: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

Several configuration options are available to customize the BDT classifier. They are summarized in Option Tables 21 and 22 and described in more detail in Sec. 8.12.2.

8.12.2 Description and implementation

Decision trees are well known classifiers that allow a straightforward interpretation as they can be visualized by a simple two-dimensional tree structure. They are in this respect similar to rectangular cuts. However, whereas a cut-based analysis is able to select only *one* hypercube as region of phase

Option	Array	Default	Predefined Values	Description
NTrees	–	200	–	Number of trees in the forest
BoostType	–	AdaBoost	AdaBoost, Bagging, RegBoost, AdaBoostR2, Grad	Boosting type for the trees in the forest
AdaBoostR2Loss	–	Quadratic	Linear, Quadratic, Exponential	Loss type used in AdaBoostR2
UseBaggedGrad	–	False	–	Use only a random subsample of all events for growing the trees in each iteration. (Only valid for GradBoost)
GradBaggingFraction	–	0.6	–	Defines the fraction of events to be used in each iteration when UseBaggedGrad=kTRUE.
Shrinkage	–	1	–	Learning rate for GradBoost algorithm
AdaBoostBeta	–	1	–	Parameter for AdaBoost algorithm
UseRandomisedTrees	–	False	–	Choose at each node splitting a random set of variables
UseNvars	–	4	–	Number of variables used if randomised tree option is chosen
UseNTrainEvent	–	N	–	Number of Training events used in each tree building if randomised tree option is chosen
UseWeightedTrees	–	True	–	Use weighted trees or simple average in classification from the forest
UseYesNoLeaf	–	True	–	Use Sig or Bkg categories, or the purity=S/(S+B) as classification of the leaf node
NodePurityLimit	–	0.5	–	In boosting/pruning, nodes with purity > NodePurityLimit are signal; background otherwise.
SeparationType	–	GiniIndex	CrossEntropy, GiniIndex, GiniIndexWithLaplace, MisClassificationError, SDivSqrtSPlusB, RegressionVariance	Separation criterion for node splitting

Option Table 21: Configuration options reference for MVA method: *BDT*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured. The table is continued in Option Table 22.

Option	Array	Default	Predefined Values	Description
`nEventsMin`	–	`max(20,NEvtsTrain/NVar`^{2}`/10)`		Minimum number of events required in a leaf node (default uses given formula)
`nCuts`	–	`20`	–	Number of steps during node cut optimisation
`PruneStrength`	–	`-1`	–	Pruning strength
`PruneMethod`	–	`CostComplexity`	`NoPruning`, `ExpectedError`, `CostComplexity`	Method used for pruning (removal) of statistically insignificant branches
`PruneBeforeBoost`	–	`False`	–	Flag to prune the tree before applying boosting algorithm
`PruningValFraction`	–	`0.5`	–	Fraction of events to use for optimizing automatic pruning.
`NNodesMax`	–	`100000`	–	Max number of nodes in tree
`MaxDepth`	–	`100000`	–	Max depth of the decision tree allowed

Option Table 22: Continuation of Option Table 21.

space, the decision tree is able to split the phase space into a large number of hypercubes, each of which is identified as either "signal-like" or "background-like", or attributed a constant event (target) value in case of a regression tree. For classification trees, the path down the tree to each leaf node represents an individual cut sequence that selects signal or background depending on the type of the leaf node.

A shortcoming of decision trees is their instability with respect to statistical fluctuations in the training sample from which the tree structure is derived. For example, if two input variables exhibit similar separation power, a fluctuation in the training sample may cause the tree growing algorithm to decide to split on one variable, while the other variable could have been selected without that fluctuation. In such a case the whole tree structure is altered below this node, possibly resulting also in a substantially different classifier response.

This problem is overcome by constructing a forest of decision trees and classifying an event on a majority vote of the classifications done by each tree in the forest. All trees in the forest are derived from the same training sample, with the events being subsequently subjected to so-called boosting (see 7), a procedure which modifies their weights in the sample. Boosting increases the statistical stability of the classifier and typically also improves the separation performance compared to a single decision tree. However, the advantage of the straightforward interpretation of the decision tree is lost. While one can of course still look at a limited number of trees trying to interpret the training result, one will hardly be able to do so for hundreds of trees in a forest. Nevertheless, the general structure of the selection can already be understood by looking at a limited number of individual trees. In many cases, the boosting performs best if applied to trees (classifiers) that, taken individually, have not much classification power, i.e. small trees.

8.12.3 Boosting, Bagging and Randomising

The different "boosting" algorithms (in the following we will call also bagging or randomised trees "boosted") available for decision trees in TMVA are currently:

- AdaBoost (see Sec. 7.1) and AdaBoostR2(see Sec. 29) for regression
- Gradient Boost (see Sec. 7.2) (not for regression)
- Bagging (see Sec. 7.3)
- Randomised Trees, like the Random Forests of L. Breiman [28]. Each tree is grown in such a way that at each split only a random subset of all variables is considered. Moreover, each tree in the forest is grown using only a (resampled) subset of the original training events. The size of the subset as well as the number of variables considered at each split can be set using the options `UseNTrainEvents` and `UseNVars`.

A possible modification of Eq. (24) for the result of the combined classifier from the forest is to use the training purity[31] in the leaf node as respective signal or background *weights* rather than relying on the binary decision. This option is chosen by setting the option `UseYesNoLeaf=False`. Such an approach however should be adopted with care as the purity in the leaf nodes is sensitive to overtraining and therefore typically overestimated. Tests performed so far with this option did not show significant performance increase. Further studies together with tree pruning are needed to better understand the behaviour of the purity-weighted BDTs.

Training (Building) a decision tree

The training, building or *growing* of a decision tree is the process that defines the splitting criteria for each node. The training starts with the root node, where an initial splitting criterion for the full training sample is determined. The split results in two subsets of training events that each go through the same algorithm of determining the next splitting iteration. This procedure is repeated until the whole tree is built. At each node, the split is determined by finding the variable and corresponding cut value that provides the best separation between signal and background. The node splitting stops once it has reached the minimum number of events which is specified in the BDT configuration (option `nEventsMin`). The leaf nodes are classified as signal or background according to the class the majority of events belongs to. If the option `UseYesNoLeaf` is set the end-nodes are classified in the same way. If `UseYesNoLeaf` is set to false the end-nodes are classified according to their purity.

A variety of separation criteria can be configured (option `SeparationType` see Option Table 22) to assess the performance of a variable and a specific cut requirement. Because a cut that selects predominantly background is as valuable as one that selects signal, the criteria are symmetric with

[31] The purity of a node is given by the ratio of signal events to all events in that node. Hence pure background nodes have zero purity.

respect to the event classes. All separation criteria have a maximum where the samples are fully mixed, i.e., at purity $p = 0.5$, and fall off to zero when the sample consists of one event class only. Tests have revealed no significant performance disparity between the following separation criteria:

- *Gini Index* [default], defined by $p \cdot (1 - p)$;
- *Cross entropy*, defined by $-p \cdot \ln(p) - (1 - p) \cdot \ln(1 - p)$;
- *Misclassification error*, defined by $1 - \max(p, 1 - p)$;
- *Statistical significance*, defined by $S/\sqrt{S + B}$;
- *Average squared error*, defined by $1/N \cdot \sum^N (y - \hat{y})^2$ for regression trees where y is the regression target of each event in the node and $\hat{y}$ is its mean value over all events in the node (which would be the estimate of y that is given by the node).

Since the splitting criterion is always a cut on a single variable, the training procedure selects *the* variable and cut value that optimises the *increase* in the separation index between the parent node and the sum of the indices of the two daughter nodes, weighted by their relative fraction of events. The cut values are optimised by scanning over the variable range with a granularity that is set via the option `nCuts`. The default value of `nCuts=20` proved to be a good compromise between computing time and step size. Finer stepping values did not increase noticeably the performance of the BDTs. However, a truly optimal cut, given the training sample, is determined by setting `nCuts=-1`. This invokes an algorithm that tests all possible cuts on the training sample and finds the best one. The latter is of course "slightly" slower than the coarse grid.

In principle, the splitting could continue until each leaf node contains only signal or only background events, which could suggest that perfect discrimination is achievable. However, such a decision tree would be strongly overtrained. To avoid overtraining a decision tree must be *pruned*.

Pruning a decision tree

Pruning is the process of cutting back a tree from the bottom up after it has been built to its maximum size. Its purpose is to remove statistically insignificant nodes and thus reduce the overtraining of the tree. It has been found to be beneficial to first grow the tree to its maximum size and then cut back, rather than interrupting the node splitting at an earlier stage. This is because apparently insignificant splits can nevertheless lead to good splits further down the tree. TMVA currently implements two tree pruning algorithms, which are set by option `PruneMethod`.

- Option `PruneMethod=ExpectedError`. For the *expected error pruning* [29] all leaf nodes for which the statistical error estimates of the parent nodes are smaller than the combined statistical error estimates of their daughter nodes are recursively deleted. The statistical error estimate of each node is calculated using the binomial error $\sqrt{p \cdot (1 - p)/N}$, where N is the number of training events in the node and p its purity. The amount of pruning is controlled by multiplying the error estimate by the fudge factor `PruneStrength`. Expected error pruning is not available for the regression trees.

- Option `PruneMethod=CostComplexity`. *Cost complexity pruning* [30] relates the number of nodes in a subtree below a node to the gain in terms of misclassified training events by the subtree compared the the node itself with *no* further splitting. The cost estimate R chosen for the misclassification of training events is given by the misclassification rate $1 - \max(p, 1-p)$ in a node. The cost complexity for this node is then defined by

$$\rho = \frac{R(\text{node}) - R(\text{subtree below that node})}{\#\text{nodes}(\text{subtree below that node}) - 1}\,. \tag{92}$$

  The node with the smallest ρ value in the tree is recursively pruned away as long as $\rho <$ `PruneStrength`. While for classification trees, one typically uses just the misclassification error in the pruning, but Gini-Index for the node splitting, regression trees use in both cases the squared error loss.

Note that the pruning is performed *after* the boosting so that the error fraction used by AdaBoost is derived from the unpruned tree.

If the `PruneStrength` option is set to a negative value, an algorithm attempts to automatically detect the optimal strength parameter. The training sample is divided into two subsamples, of which only one is used for training, while the other one serves for validation. The tree is pruned sequentially starting from the node which has the smallest value of the cost-complexity in the tree. After each pruning step the performance of the tree is assessed using the validation sample. This process is repeated until the ROOT node would be pruned. As optimal prune strength for this tree the value is chose which corresponds to the best performing tree using the validation sample.

While this type of pruning obviously gives the "optimally pruned tree" given the training data, it is not completely clear yet if this also applies for the tree in the forest. Currently it looks as if in TMVA, better results for the whole forest are often achieved when pruning is not applied, but rather the maximal tree depth is set to a relatively small value (3 or 4) already during the tree building phase.

Note that the Gradient boost does not apply a pruning algorithm and ignores option `PruneMethod`. In this case it is recommended that the user restricts the number of nodes in the tree to values between 5 to 20 by using option `NNodesMax` or the maximal allowed depth of the tree `MaxDepth`.

8.12.4 Variable ranking

A ranking of the BDT input variables is derived by counting how often the variables are used to split decision tree nodes, and by weighting each split occurrence by the separation gain-squared it has achieved and by the number of events in the node [30]. This measure of the variable importance can be used for a single decision tree as well as for a forest.

8.12.5 Performance

Only limited experience has been gained so far with boosted decision trees in HEP. In the literature decision trees are sometimes referred to as the best "out of the box" classifiers. This is because little tuning is required in order to obtain reasonably good results. This is due to the simplicity of the method where each training step (node splitting) involves only a one-dimensional cut optimisation. Decision trees are also insensitive to the inclusion of poorly discriminating input variables. While for artificial neural networks it is typically more difficult to deal with such additional variables, the decision tree training algorithm will basically ignore non-discriminating variables as for each node splitting only the best discriminating variable is used. However, the simplicity of decision trees has the drawback that their theoretically best performance on a given problem is generally inferior to other techniques like neural networks. This is seen for example using the academic training samples included in the TMVA package. For this sample, which has equal RMS but shifted mean values for signal and background and linear correlations between the variables only, the Fisher discriminant provides theoretically optimal discrimination results. While the artificial neural networks are able to reproduce this optimal selection performance the BDTs always fall short in doing so. However, in other academic examples with more complex correlations or real life examples, the BDTs often outperform the other techniques. This is because either there are not enough training events available that would be needed by the other classifiers, or the optimal configuration (i.e. how many hidden layers, which variables) of the neural network has not been specified. We have only very limited experience at the time with the regression, hence cannot really comment on the performance in this case.

8.13 Predictive learning via rule ensembles (RuleFit)

This classifier is a TMVA implementation of Friedman-Popescu's RuleFit method described in [32]. Its idea is to use an ensemble of so-called *rules* to create a scoring function with good classification power. Each rule r_i is defined by a sequence of cuts, such as

$$\begin{aligned} r_1(\mathbf{x}) &= I(x_2 < 100.0) \cdot I(x_3 > 35.0)\,, \\ r_2(\mathbf{x}) &= I(0.45 < x_4 < 1.00) \cdot I(x_1 > 150.0)\,, \\ r_3(\mathbf{x}) &= I(x_3 < 11.00)\,, \end{aligned}$$

where the x_i are discriminating input variables, and $I(\dots)$ returns the truth of its argument. A rule applied on a given event is non-zero only if all of its cuts are satisfied, in which case the rule returns 1.

The easiest way to create an ensemble of rules is to extract it from a forest of decision trees (cf. Sec. 8.12). Every node in a tree (except the root node) corresponds to a sequence of cuts required to reach the node from the root node, and can be regarded as a rule. Hence for the tree illustrated in Fig. 18 on page 105 a total of 8 rules can be formed. Linear combinations of the rules in the ensemble are created with coefficients (rule weights) calculated using a regularised minimisation procedure [33]. The resulting linear combination of all rules defines a *score* function (see below) which provides the RuleFit response $y_{\mathrm{RF}}(\mathbf{x})$.

In some cases a very large rule ensemble is required to obtain a competitive discrimination between signal and background. A particularly difficult situation is when the true (but unknown) scoring function is described by a linear combination of the input variables. In such cases, e.g., a Fisher discriminant would perform well. To ease the rule optimisation task, a linear combination of the input variables is added to the model. The minimisation procedure will then select the appropriate coefficients for the rules *and* the linear terms. More details are given in Sec. 8.13.2 below.

8.13.1 Booking options

The RuleFit classifier is booked via the command:

```
factory->BookMethod( Types::kRuleFit, "RuleFit", "<options>" );
```

Code Example 47: Booking of RuleFit: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The RuleFit configuration options are given in Option Table 23.

8.13.2 Description and implementation

As for all TMVA classifiers, the goal of the rule learning is to find a classification function $y_{\mathrm{RF}}(\mathbf{x})$ that optimally classifies an event according to the tuple of input observations (variables) $\mathbf{x}$. The classification function is written as

$$y_{\mathrm{RF}}(\mathbf{x}) = a_0 + \sum_{m=1}^{M_R} a_m f_m(\mathbf{x}) \,, \tag{93}$$

where the set $\{f_m(\mathbf{x})\}_{M_R}$ forms an ensemble of *base learners* with M_R elements. A base learner may be any discriminating function derived from the training data. In our case, they consist of rules and linear terms as described in the introduction. The complete model then reads

$$y_{\mathrm{RF}}(\mathbf{x}) = a_0 + \sum_{m=1}^{M_R} a_m r_m(\mathbf{x}) + \sum_{i=1}^{n_{\mathrm{var}}} b_i x_i \,. \tag{94}$$

To protect against outliers, the variables in the linear terms are modified to

$$x'_i = \min(\delta_i^+, \max(\delta_i^-)) \,, \tag{95}$$

where $\delta_i^{\pm}$ are the lower and upper β quantiles[32] of the variable x_i. The value of β is set by the option `LinQuantile`. If the variables are used "as is", they may have an unequal *a priori* influence

[32] Quantiles are points taken at regular intervals from the PDF of a random variable. For example, the 0.5 quantile corresponds to the median of the PDF.

Option	Array	Default	Predefined Values	Description
GDTau	–	-1	–	Gradient-directed (GD) path: default fit cut-off
GDTauPrec	–	0.01	–	GD path: precision of tau
GDStep	–	0.01	–	GD path: step size
GDNSteps	–	10000	–	GD path: number of steps
GDErrScale	–	1.1	–	Stop scan when error > scale*errmin
LinQuantile	–	0.025	–	Quantile of linear terms (removes outliers)
GDPathEveFrac	–	0.5	–	Fraction of events used for the path search
GDValidEveFrac	–	0.5	–	Fraction of events used for the validation
fEventsMin	–	0.1	–	Minimum fraction of events in a splittable node
fEventsMax	–	0.9	–	Maximum fraction of events in a splittable node
nTrees	–	20	–	Number of trees in forest.
ForestType	–	AdaBoost	AdaBoost, Random	Method to use for forest generation
RuleMinDist	–	0.001	–	Minimum distance between rules
MinImp	–	0.01	–	Minimum rule importance accepted
Model	–	ModRuleLinear	ModRule, ModRuleLinear, ModLinear	Model to be used
RuleFitModule	–	RFTMVA	RFTMVA, RFFriedman	Which RuleFit module to use
RFWorkDir	–	./rulefit	–	Friedman's RuleFit module (RFF): working dir
RFNrules	–	2000	–	RFF: Mximum number of rules
RFNendnodes	–	4	–	RFF: Average number of end nodes

Option Table 23: Configuration options reference for MVA method: *RuleFit*. Values given are defaults. If predefined categories exist, the default category is marked by a '⋆'. The options in Option Table 9 on page 58 can also be configured.

relative to the rules. To counter this effect, the variables are normalised

$$x'_i \to \sigma_r \cdot x'_i/\sigma_i \,, \tag{96}$$

where σ_r and σ_i are the estimated standard deviations of an ensemble of rules and the variable x'_i, respectively.

Rule generation

The rules are extracted from a forest of decision trees. There are several ways to generate a forest. In the current RuleFit implementation, each tree is generated using a fraction of the training sample. The fraction depends on which method is used for generating the forest. Currently two methods are supported (selected by option `ForestType`); *AdaBoost* and *Random Forest*. The first method is described in Sec. 8.12.2. In that case, the whole training set is used for all trees. The diversity is obtained through using different event weights for each tree. For a random forest, though, the diversity is created by training each tree using random sub-samples. If this method is chosen, the fraction is calculated from the training sample size N (signal and background) using the empirical formula [34]

$$f = \min(0.5, (100.0 + 6.0 \cdot \sqrt{N})/N)\,. \tag{97}$$

By default, `AdaBoost` is used for creation of the forest. In general it seems to perform better than the random forest.

The topology of each tree is controlled by the parameters `fEventsMin` and `fEventsMax`. They define a range of fractions which are used to calculate the minimum number of events required in a node for further splitting. For each tree, a fraction is drawn from a uniform distribution within the given range. The obtained fraction is then multiplied with the number of training events used for the tree, giving the minimum number of events in a node to allow for splitting. In this way both large trees (small fraction) giving complex rules and small trees (large fraction) for simple rules are created. For a given forest of N_t trees, where each tree has n_ℓ leaf nodes, the maximum number of possible rules is

$$M_{R,\max} = \sum_{i=1}^{N_t} 2(n_{\ell,i} - 1)\,. \tag{98}$$

To prune similar rules, a *distance* is defined between two *topologically equal* rules. Two rules are topologically equal if their cut sequences follow the same variables only differing in their cut values. The rule distance used in TMVA is then defined by

$$\delta_R^2 = \sum_i \frac{\delta_{i,L}^2 + \delta_{i,U}^2}{\sigma_i^2}\,, \tag{99}$$

where $\delta_{i,L(U)}$ is the difference in lower (upper) limit between the two cuts containing the variable x_i, $i = 1, \ldots, n_{\text{var}}$. The difference is normalised to the RMS-squared σ_i^2 of the variable. Similar rules with a distance smaller than `RuleMinDist` are removed from the rule ensemble. The parameter can be tuned to improve speed and to suppress noise. In principle, this should be achieved in the fitting procedure. However, pruning the rule ensemble using a distance cut will reduce the fitting time and will probably also reduce the number of rules in the final model. Note that the cut should be used with care since a too large cut value will deplete the rule ensemble and weaken its classification performance.

Fitting

Once the rules are defined, the coefficients in Eq. (94) are fitted using the training data. For details, the fitting method is described in [33]. A brief description is provided below to motivate the corresponding RuleFit options.

A *loss function* $L(y_{\mathrm{RF}}(\mathbf{x})|\hat{y})$, given by the "squared-error ramp" [33]

$$L(y_{\mathrm{RF}}|\hat{y}) = (\hat{y} - H(y_{\mathrm{RF}}))^2 \,, \tag{100}$$

where $H(y) = \max(-1, \min(y_{\mathrm{RF}}, 1))$, quantifies the "cost" of misclassifying an event of given true class $\hat{y}$. The *risk* R is defined by the expectation value of L with respect to $\mathbf{x}$ and the true class. Since the true distributions are generally not known, the average of N training events is used as an estimate

$$R = \frac{1}{N} \sum_{i=1}^{N} L(y_{\mathrm{RF}}(\mathbf{x}_i)|\hat{y}_i) \,. \tag{101}$$

A line element in the parameter space of the rule weights (given by the vector $\mathbf{a}$ of all coefficients) is then defined by

$$\mathbf{a}(\epsilon + \delta\epsilon) = \mathbf{a}(\epsilon) + \delta\epsilon \cdot \mathbf{g}(\epsilon) \,, \tag{102}$$

where $\delta\epsilon$ is a positive small increment and $\mathbf{g}(\epsilon)$ is the negative derivative of the estimated risk R, evaluated at $\mathbf{a}(\epsilon)$. The estimated risk-gradient is evaluated using a sub-sample (`GDPathEveFrac`) of the training events.

Starting with all weights set to zero, the consecutive application of Eq. (102) creates a path in the $\mathbf{a}$ space. At each step, the procedure selects only the gradients g_k with absolute values greater than a certain fraction (τ) of the largest gradient. The fraction τ is an *a priori* unknown quantity between 0 and 1. With $\tau = 0$ all gradients will be used at each step, while only the strongest gradient is selected for $\tau = 1$. A measure of the "error" at each step is calculated by evaluating the risk (Eq. 101) using the validation sub-sample (`GDValidEveFrac`). By construction, the risk will always decrease at each step. However, for the validation sample the value will increase once the model starts to be overtrained. Currently, the fitting is crudely stopped when the error measure is larger than `GDErrScale` times the minimum error found. The number of steps is controlled by `GDNSteps` and the step size ($\delta\epsilon$ in Eq. 102) by `GDStep`.

If the selected τ (`GDTau`) is a negative number, the best value is estimated by means of a scan. In such a case several paths are fitted in parallel, each with a different value of τ. The number of paths created depend on the required precision on τ given by `GDTauPrec`. By only selecting the paths being "close enough" to the minimum at each step, the speed for the scan is kept down. The path leading to the lowest estimated error is then selected. Once the best τ is found, the fitting proceeds until a minimum is found. A simple example with a few scan points is illustrated in Fig. 19.

8.13.3 Variable ranking

Since the input variables are normalised, the ranking of variables follows naturally from the coefficients of the model. To each rule m ($m = 1, \ldots, M_R$) can be assigned an importance defined

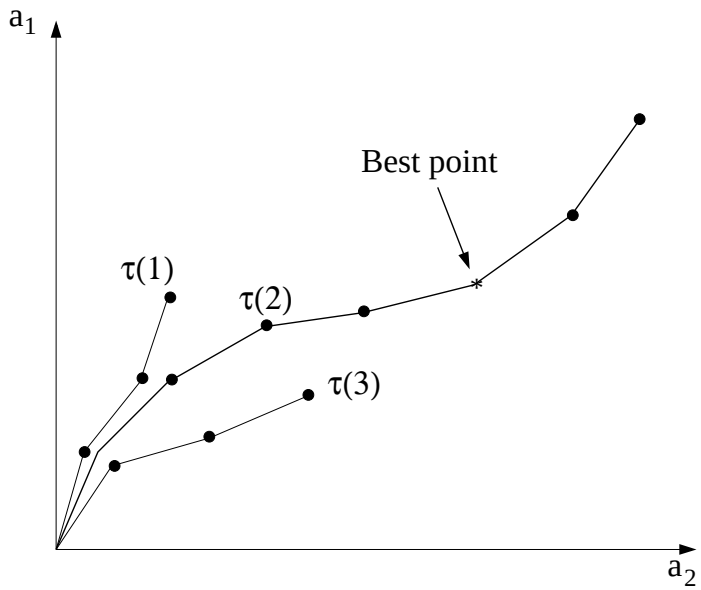

Figure 19: An example of a path scan in two dimensions. Each point represents an ϵ in Eq. (102) and each step is given by $\delta\epsilon$. The direction along the path at each point is given by the vector $\mathbf{g}$. For the first few points, the paths $\tau(1,2,3)$ are created with different values of τ. After a given number of steps, the best path is chosen and the search is continued. It stops when the best point is found. That is, when the estimated error-rate is minimum.

by

$$I_m = |a_m|\sqrt{s_m(1.0 - s_m)}\,, \tag{103}$$

where s_m is the *support* of the rule with the following definition

$$s_m = \frac{1}{N}\sum_{n=1}^{N} r_m(\mathbf{x}_n)\,. \tag{104}$$

The support is thus the average response for a given rule on the data sample. A large support implies that many events pass the cuts of the rule. Hence, such rules cannot have strong discriminating power. On the other hand, rules with small support only accept few events. They may be important for these few events they accept, but they are not in the overall picture. The definition (103) for the rule importance suppresses rules with both large and small support.

For the linear terms, the definition of importance is

$$I_i = |b_i| \cdot \sigma_i\,, \tag{105}$$

so that variables with small overall variation will be assigned a small importance.

A measure of the variable importance may then be defined by

$$J_i = I_i + \sum_{m|x_i \in r_m} I_m/q_m\,, \tag{106}$$

where the sum is over all rules containing the variable x_i, and q_m is the number of variables used in the rule r_m. This is introduced in order to share the importance equally between all variables in rules with more than one variable.

8.13.4 Friedman's module

By setting `RuleFitModule` to `RFFriedman`, the interface to Friedman's RuleFit is selected. To use this module, a separate setup is required. If the module is selected in a run prior to setting up the environment, TMVA will stop and give instructions on how to proceed. A command sequence to setup Friedman's RuleFit in a UNIX environment is:

```
~> mkdir rulefit
~> cd rulefit
~> wget http://www-stat.stanford.edu/~jhf/r-rulefit/linux/rf_go.exe
~> chmod +x rf_go.exe
```

Code Example 48: The first line creates a working directory for Friedman's module. In the third line, the binary executable is fetched from the official web-site. Finally, it is made sure that the module is executable.

As of this writing, binaries exists only for Linux and Windows. Check J. Friedman's home page at http://www-stat.stanford.edu/~jhf for updated information. When running this module from TMVA, make sure that the option `RFWorkDir` is set to the proper working directory (default is `./rulefit`). Also note that only the following options are used: `Model`, `RFWorkDir`, `RFNrules`, `RFNendnodes`, `GDNSteps`, `GDStep` and `GDErrScale`. The options `RFNrules` and `RFNendnodes` correspond in the package by Friedman-Popescu to the options `max.rules` and `tree.size`, respectively. For more details, the reader is referred to Friedman's RuleFit manual [34].

Technical note

The module `rf_go.exe` communicates with the user by means of both ASCII and binary files. This makes the input/output from the module machine dependant. TMVA reads the output from `rf_go.exe` and produces the normal machine independent weight (or class) file. This can then be used in other applications and environments.

8.13.5 Performance

Rule ensemble based learning machines are not yet well known within the HEP community, although they start to receive some attention [35]. Apart from RuleFit [32] other rule ensemble learners exists, such as SLIPPER [36].

The TMVA implementation of RuleFit follows closely the original design described in Ref. [32].

Currently the performance is however slightly less robust than the one of the Friedman-Popescu package. Also, the experience using the method is still scarce at the time of this writing.

To optimise the performance of RuleFit several strategies can be employed. The training consists of two steps, rule generation and rule ensemble fitting. One approach is to modify the complexity of the generated rule ensemble by changing either the number of trees in the forest, or the complexity of each tree. In general, large tree ensembles with varying trees sizes perform better than short non-complex ones. The drawback is of course that fitting becomes slow. However, if the fitting performs well, it is likely that a large amount of rules will have small or zero coefficients. These can be removed, thus simplifying the ensemble. The fitting performance can be improved by increasing the number of steps along with using a smaller step size. Again, this will be at the cost of speed performance although only at the training stage. The setting for the parameter τ may greatly affect the result. Currently an automatic scan is performed by default. In general, it should find the optimum τ. If in doubt , the user may set the value explicitly. In any case, the user is initially advised to use the automatic scan option to derive the best path.

9 Combining MVA Methods

In challenging classification or regression problems with a high demand for optimisation, or when treating feature spaces with strongly varying properties, it can be useful to combined MVA methods. There is large room for creativity inherent in such combinations. For TMVA we distinguish three classes of combinations:

1. *Boosting* MVA methods,
2. *Categorising* MVA methods,
3. Building *committees* of MVA methods.

While the first two combined methods are available since TMVA 4.1, the third one is still under development. The MVA *booster* is a generalisation of the Boosted Decision Trees approach (see Sec. 8.12) to any method in TMVA. *Category methods* allow the user to specify sub-regions of phase space, which are assigned by requirements on input or spectator variables, and which define disjoint sub-populations of the training sample, thus improving the modelling of the training data. Finally, *Committee methods* allow one to input MVA methods into other MVA methods, a procedure that can be arbitrarily chained.

All of these combined methods are of course MVA methods themselves, treated just like any other method in TMVA for training, testing, evaluation and application. This also allows to categorise a committee method, for example.

9.1 Boosted classifiers

Since generalised boosting is not yet available for regression in TMVA, we restrict the following discussion to classification applications. A boosted classifier is a combination of a collection of classifiers of the same type trained on the same sample but with different events weights.[33] The response of the final classifier is a weighted response of each individual classifier in the collection. The boosted classifier is potentially more powerful and more stable with respect to statistical fluctuations in the training sample. The latter is particularly the case for bagging as "boost" algorithm (cf. Sec. 7.3, page 55).

The following sections do not apply to decision trees. We refer to Sec. 8.12 (page 104) for a description of boosted decision trees. In the current version of TMVA only the AdaBoost and Bagging algorithms are implemented for the boost of arbitrary classifiers. The boost algorithms are described in detail in Sec. 7 on page 52.

9.1.1 Booking options

To book a boosted classifier, one needs to add the booster options to the regular classifier's option string. The minimal option required is the number of boost iterations `Boost_Num`, which must be set to a value larger than zero. Once the Factory detects a `Boost_Num`⟩`0` in the option string it books a boosted classifier and passes all boost options (recognised by the prefix `Boost_`) to the Boost method and the other options to the boosted classifier.

```
factory->BookMethod( TMVA::Types::kLikelihood, "BoostedLikelihood",
       "Boost_Num=10:Boost_Type=Bagging:Spline=2:NSmooth=5:NAvEvtPerBin=50" );
```

Code Example 49: Booking of the boosted classifier: the first argument is the predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. All options with the prefix `Boost_` (in this example the first two options) are passed on to the boost method, the other options are provided to the regular classifier (which in this case is Likelihood). Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The boost configuration options are given in Option Table 24.

The options most relevant for the boost process are the number of boost iterations, `Boost_Num`, and the choice of the boost algorithm, `Boost_Type`. In case of `Boost_Type=AdaBoost`, the option `Boost_Num` describes the maximum number of boosts. The algorithm is iterated until an error rate of 0.5 is reached or until `Boost_Num` iterations occurred. If the algorithm terminates after to few iterations, the number might be extended by decreasing the β variable (option `Boost_AdaBoostBeta`). Within the AdaBoost algorithm a decision must be made how to classify an event, a task usually done by the user. For some classifiers it is straightforward to set a cut on the MVA response to define signal-like events. For the others, the MVA cut is chosen that the error rate is minimised. The option

[33] The Boost method is at the moment only applicable to classification problems.

Option	Array	Default	Predefined Values	Description
Boost_Num	–	100	–	Number of times the classifier is boosted
Boost_MonitorMethod	–	True	–	Whether to write monitoring histogram for each boosted classifier
Boost_Type	–	AdaBoost	AdaBoost, Bagging	Boosting type for the classifiers
Boost_MethodWeightType	–	ByError	ByError, Average, LastMethod	How to set the final weight of the boosted classifiers
Boost_RecalculateMVACut	–	True	–	Whether to recalculate the classifier MVA Signallike cut at every boost iteration
Boost_AdaBoostBeta	–	1	–	The ADA boost parameter that sets the effect of every boost step on the events' weights
Boost_Transform	–	step	step, linear, log	Type of transform applied to every boosted method linear, log, step

Option Table 24: Boosting configuration options. These options can be simply added to a simple classifier's option string or used to form the option string of an explicitly booked boosted classifier.

`Boost_RecalculateMVACut` determines whether this cut should be recomputed for every boosting iteration. In case of Bagging as boosting algorithm the number of boosting iterations always reaches `Boost_Num`.

By default boosted classifiers are combined as a weighted average with weights computed from the misclassification error (option `Boost_MethodWeightType=ByError`). It is also possible to use the arithmetic average instead (`Boost_MethodWeightType=Average`).

9.1.2 Boostable classifiers

The boosting process was originally introduced for simple classifiers. The most commonly boosted classifier is the decision tree (DT – cf. Sec. 8.12, page 104). Decision trees need to be boosted a few hundred times to effectively stabilise the BDT response and achieve optimal performance.

Another simple classifier in the TMVA package is the Fisher discriminant (cf. Sec. 8.7, page 84 – which is equivalent to the linear discriminant described in Sec. 8.8). Because the output of a Fisher discriminant represents a linear combination of the input variables, a linear combination of different Fisher discriminants is again a Fisher discriminant. Hence linear boosting cannot improve the performance. It is nevertheless possible to effectively boost a linear discriminant by applying the linear combination not on the discriminant's output, but on the actual classification results provided.[34] This corresponds to a "non-linear" transformation of the Fisher discriminant output according to a

[34] Note that in the TMVA standard example, which uses linearly correlated, Gaussian-distributed input variables

step function. The Boost method in TMVA also features a fully non-linear transformation that is directly applied to the classifier response value. Overall, the following transformations are available:

- *linear:* no transformation is applied to the MVA output,
- *step:* the output is -1 below the step and $+1$ above (default setting),
- *log:* logarithmic transformation of the output.

The macro `Boost.C` (residing in the `macros` (`test`) directory for the sourceforge (ROOT) version of TMVA) provides examples for the use of these transformations to boost a Fisher discriminant. We point out that the performance of a boosted classifier strongly depends on its characteristics as well as on the nature of the input data. A careful adjustment of options is required if AdaBoost is applied to an arbitrary classifier, since otherwise it might even lead to a worse performance than for the unboosted method.

9.1.3 Monitoring tools

The current GUI provides figures to monitor the boosting process. Plotted are the boost weights, the classifier weights in the boost ensemble, the classifier error rates, and the classifier error rates using unboosted event weights. In addition, when the option `Boost_MonitorMethod=T` is set, monitoring histograms are created for each classifier in the boost ensemble. The histograms generated during the boosting process provide useful insight into the behaviour of the boosted classifiers and help to adjust to the optimal number of boost iterations. These histograms are saved in a separate folder in the output file, within the folder of `MethodBoost/<Title>/`. Besides the specific classifier monitoring histograms, this folder also contains the MVA response of the classifier for the training and testing samples.

9.1.4 Variable ranking

The present boosted classifier implementation does not provide a ranking of the input variables.

9.2 Category Classifier

The *Category method* allows the user to separate the training data (and accordingly the application data) into disjoint sub-populations exhibiting different properties. The separation into phase space regions is done by applying requirements on the input and/or *spectator* variables (variables defined as spectators are not used as direct inputs to MVA methods). It thus reaches beyond a simple sub-classification of the original feature space. In each of these disjoint regions (each event must belong to one and only one region), an independent training is performed using the most appropriate MVA

for signal and background, a single Fisher discriminant already provides the theoretically maximum separation power. Hence on this example, no further gain can be expected by boosting, no matter what "tricks" are applied.

method, training options and set of training variables in that zone. The division into categories in presence of distinct sub-populations reduces the correlations between the training variables, improves the modelling, and hence increases the classification and regression performance.[35] Another common application is when variables are not available in some phase space regions. For example, a sub-detector may only cover a fraction of the training data sample. By defining two categories the variables provided by this detector could be excluded from the training in the phase space region outside of its fiducial volume. Finally, an advantage of categorisation, often exploited in likelihood fits, also lies in the gain of signal significance when the fraction of signal events in the event sample differs between the categories.

We point out that the explicit use of the Category method is not mandatory. The categorisation functionality can be achieved by hand by training independent MVAs in each of the disjoint regions. Our experience however shows that often the corresponding bookkeeping is considered too cumbersome and the required categorisation is not performed, thus leading to performance loss.[36] The Category method is straightforward to implement and should help in almost all analyses to better describe the training sample. Performance comparison with the corresponding non-categorised method should guide the user in his/her choice.

For the current release, the Category method is only implemented for classification. Extension to regression is however straightforward and should follow soon.

9.2.1 Booking options

The Category method is booked via the command:

```
TMVA::IMethod* category = factory->BookMethod( TMVA::Types::kCategory,
                                               "Category",
                                               "<options>" );
```

Code Example 50: Booking of the Category method: the first argument is a predefined enumerator, the second argument is a user-defined string identifier, and the third argument is the configuration options string. Individual options are separated by a ':'. See Sec. 3.1.5 for more information on the booking.

The categories and sub-classifiers are defined as follows:

[35] This is particularly the case for projective likelihood methods, even when including prior decorrelation of the input data (cf. Sec. 8.2).

[36] In the above example with the reduced detector fiducial acceptance, it often occurs that non-physical values ("−99") are assigned to the variables for events where they are not available. Inserting such meaningless (apart from indicating the category) values into an MVA penalises its performance.

```
TMVA::MethodCategory* mcategory = dynamic_cast<TMVA::MethodCategory*>(category);
mcategory->AddCategory( "<cut>",
                        "<variables>",
                        TMVA::Types::<enumerator>,
                        "<sub-classifier name>",
                        "<sub-classifier options>" );
```

Code Example 51: Adding a category to the Category method: the first argument is the cut that defines the category, the second defines is the set of variables used to train this sub-classifier, the third argument is the predefined enumerator of the sub-classifier, the fourth argument is a user-defined string identifier of the sub-classifier, and the last argument sets the configuration options of the sub-classifier. Individual variables and options are both separated by a ':'. See Sec. 9.2.2 for further information on cuts and variables. The dynamic cast is required to allow access to specific Category members.

There are no options implemented for the Category method itself, apart from those common to all methods (cf. Option Table 9 on page 58). The available options for the sub-classifiers are found in the corresponding sections of this Users Guide.

The following example illustrates how sub-classifiers are added:

```
mcategory->AddCategory( "abs(eta)<=1.3",
                        "var1:var2:var3:var4:",
                        TMVA::Types::kFisher,
                        "Category_Fisher",
                        "H:!V:Fisher" );
mcategory->AddCategory( "abs(eta)>1.3",
                        "var1:var2:var3:",
                        TMVA::Types::kFisher,
                        "Category_Fisher",
                        "H:!V:Fisher" );
```

Code Example 52: Adding a category with sub-classifier of type Fisher: the cut in the first argument defines the region to which each of the sub-classifiers are applied, the second argument is the list of variables which are to be considered for this region (in this example `var4` is not available when `abs(eta)>1.3`), the third argument contains the internal enumerator for a classifier of type Fisher (the MVA method may differ between categories, the fourth argument is the user-defined identifier, and the fifth argument contains the options string for the Fisher classifier. The variable `eta` must have been defined as an input or spectator variable.

9.2.2 Description and implementation

The Category method is implemented in an entirely transparent way, conserving full flexibility for the user on which input variables, MVA method and options to use in each category, and how the

category should be defined.

The categories are defined via cuts on the input and/or spectator variables. These cuts can be arbitrary boolean expressions or may involve any number of variables. In case of particularly intricate category definitions it may turn out advantageous for the user to pre-define and pre-fill category variables, and use those as spectators in TMVA to set an event's category. It is required that the cuts are defined in such a way that they create disjoint event samples. TMVA will issue a fatal error in case of an ambiguous category definition.

The variables used for the training of each of the sub-classifiers can be chosen from the variables and spectators that have been defined in the TMVA script. The individual variables have to be separated by a ':'. The distinction between input variables and spectators is necessary to allow performance comparisons between categorised and non-categorised methods within the same training and evaluation job (spectator variables would not be used by non-categorised methods).

The advantage offered by the Category classifier is illustrated with the following academic example. We consider four uncorrelated and Gaussian distributed input variables for training. The mean values of the variables are shifted between signal and background, hence providing classification information. In addition, all mean values (signal *and* background) depend on a spectator variable `eta`: for `abs(eta)>1.3` (`abs(eta)<=1.3`) they have a positive (negative) offset. Hence, the Gaussians are broader and the variables receive an effective correlation when integrating over `eta`. The distributions for signal and background of one out of the four variables is shown for both `abs(eta)` categories in Fig. 20. Figure 21 shows the ROC curves (cf. Sec. 2.7) of the categorised and non-categorised Fisher and Projective Likelihood classifiers. Optimum performance is recovered by the Category classifier.

9.2.3 Variable ranking

Due to the nature of the Category classifier it does not provide a general variable ranking. As far as available, the sub-classifiers will provide ranking for each category.

9.2.4 Performance

The performance of the Category method depends entirely on the judicious choice of the categories, properly separating events with different properties, and on the performance of the sub-methods within these categories. Control by performance comparison with the corresponding non-categorised methods allows to measure the performance gain, and to assess whether categorisation improves the classification.

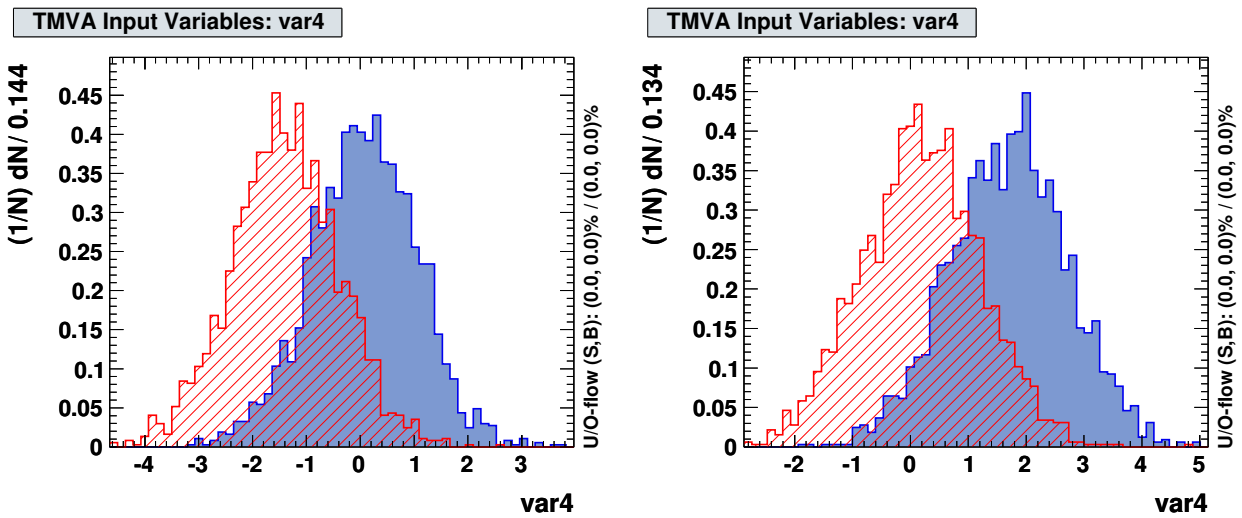

Figure 20: Examples for the usage of categories. The variable `var4` depends on `abs(eta)`: for `abs(eta)<=1.3` (left plot), the signal and background Gaussian distributions are both shifted to lower values, while for `abs(eta)>1.3` (right plot) shifts to larger values have been applied. Ignoring this dependence, broadens the inclusive Gaussian distributions and generates variable correlations (here a coefficient of +0.4 between `var3` and `var4`) due to the existence of distinct subsamples. The Category classifier treats the two cases independently, thus recovering optimal separation.

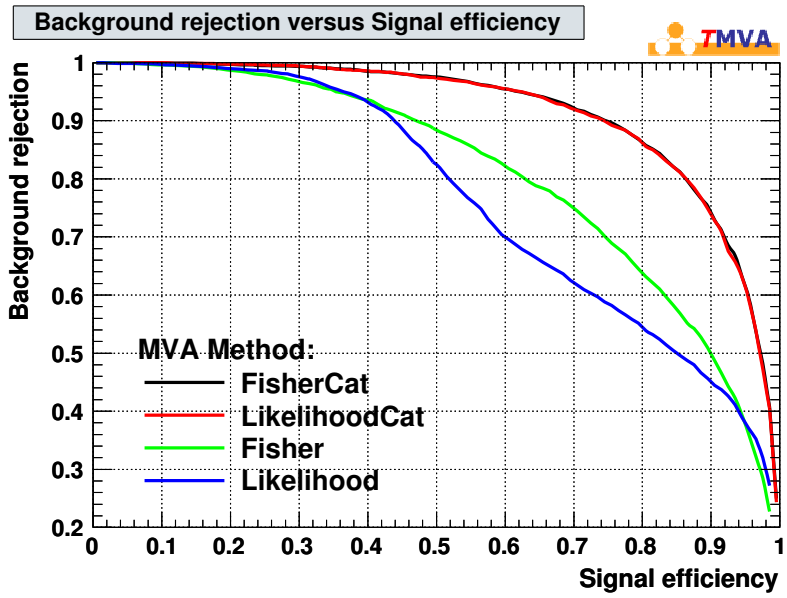

Figure 21: ROC curves for categorised and non-categorised Fisher and Projective Likelihood classifiers applied to an academic four-Gaussian example with `eta`-dependent mean values distinguishing two categories (cf. Fig. 20 for the category-dependent signal and background distributions of one of the four input variables). The categorised ("...Cat") classifiers recover optimum performance (their red and black curves are superimposed).

CRITERIA		MVA METHOD									
		Cuts	Likelihood	PDE-RS / k-NN	PDE-Foam	H-Matrix	Fisher / LD	MLP	BDT	Rule-Fit	SVM
Performance	No or linear correlations	⋆	⋆⋆	⋆	⋆	⋆	⋆⋆	⋆⋆	⋆	⋆⋆	⋆
	Nonlinear correlations	○	○	⋆⋆	⋆⋆	○	○	⋆⋆	⋆⋆	⋆⋆	⋆⋆
Speed	Training	○	⋆⋆	⋆⋆	⋆⋆	⋆⋆	⋆⋆	⋆	○	⋆	○
	Response	⋆⋆	⋆⋆	○	⋆	⋆⋆	⋆⋆	⋆⋆	⋆	⋆⋆	⋆
Robustness	Overtraining	⋆⋆	⋆	⋆	⋆	⋆⋆	⋆⋆	⋆	○	⋆	⋆⋆
	Weak variables	⋆⋆	⋆	○	○	⋆⋆	⋆⋆	⋆	⋆⋆	⋆	⋆
Curse of dimensionality		○	⋆⋆	○	○	⋆⋆	⋆⋆	⋆	⋆	⋆	
Transparency		⋆⋆	⋆⋆	⋆	⋆	⋆⋆	⋆⋆	○	○	○	○

Table 6: Assessment of MVA method properties. The symbols stand for the attributes "good" (⋆⋆), "fair" (⋆) and "bad" (○). "Curse of dimensionality" refers to the "burden" of required increase in training statistics and processing time when adding more input variables. See also comments in the text. The FDA method is not listed here since its properties depend on the chosen function.

10 Which MVA method should I use for my problem?

There is obviously no generally valid answer to this question. To guide the user, we have attempted a coarse assessment of various MVA properties in Table 6. Simplicity is a virtue, but only if it is not at the expense of significant loss of discrimination power. Robustness with respect to overtraining could become an issue when the training sample is scarce. Some methods require more attention than others in this regard. For example, boosted decision trees are particularly vulnerable to overtraining if used without care.[37] To circumvent overtraining a problem-specific adjustment of the pruning strength parameter is required.

To assess whether a linear discriminant analysis (LDA) could be sufficient for a classification (regression) problem, the user is advised to analyse the correlations among the discriminating variables (among the variables and regression target) by inspecting scatter and profile plots (it is not enough to print the correlation coefficients, which by definition are linear only). Using an LDA greatly reduces the number of parameters to be adjusted and hence allow smaller training samples. It usually is robust with respect to generalisation to larger data samples. For moderately intricate problems, the function discriminant analysis (FDA) with some added nonlinearity may be found sufficient. It is always useful to cross-check its performance against several of the sophisticated nonlinear methods to see how much can be gained over the use of the simple and very transparent FDA.

[37] However, experience shows that the BDT performance is amazingly robust – even for strongly overtrained decision trees.

For problems that require a high degree of optimisation and allow to use a large number of input variables, complex nonlinear methods like neural networks, the support vector machine, boosted decision trees and/or RuleFit are more appropriate.

Very involved multi-dimensional variable correlations with strong nonlinearities are usually best mapped by the multidimensional probability density estimators such as PDE-RS, k-NN and PDE-Foam, requiring however a reasonably low number of input variables.

For RuleFit classification we emphasise that the TMVA implementation differs from Friedman-Popescu's original code [32], with slightly better robustness and out-of-the-box performance for the latter version. In particular, the behaviour of the original code with respect to nonlinear correlations and the curse of dimensionality would have merited two stars.[38] We also point out that the excellent performance for by majority linearly correlated input variables is achieved somewhat artificially by adding a Fisher-like term to the RuleFit classifier (this is the case for both implementations, cf. Sec. 8.13 on page 111).

In Fig 5 (page 12) we have shown the example of ROC-curves obtained from various linear and non-linear classifiers on a simple toy Monte Carlo sample of linear correlated Gaussian distribution functions. Here all classifiers capable of dealing with such type of correlations are equally performing. Only the projective likelihood method, which ignores all correlations, performs significantly worse (it recovers optimum performance after prior decorrelation of the variables). For such a problem, a Fisher (aka linear) discriminant is optimal by construction, and the non-linear methods effectively reduce in the training to linear discriminants achieving competitive, but somewhat less robust performance.

For another academic toy Monte Carlo example however, which exhibits strongly non-linear correlations in form of a "Schachbrett" (chess board with two-dimensional Gaussian-distributed signal and background spots – see Fig. 22), the limitations of the linear classifiers, being unable to capture the features of the data, become manifest. Non-linear classifiers, once appropriately optimised such as choosing a sufficiently complex network architecture for the MLP, or the proper kernel functions in the SVM, all non-linear methods reproduce or approximate the theoretical limit of this classification problem (indicated by the thick red line in the ROCc curve in the lower plot of Fig. 22).

11 TMVA implementation status summary for classification and regression

All TMVA methods are fully operational for user analysis, requiring training, testing, evaluating and reading for the application to unknown data samples. Additional features are optional and – despite our attempts to provide a fully transparent analysis – not yet uniformly available. A status summary is given in Table 7 and annotated below.

Although since TMVA 4 the framework supports multi-dimensional MVA outputs it has not yet been implemented for classification. For regression, only a few methods are fully multi-target capable so

[38] An interface to Friedman-Popescu's original code can be requested from the TMVA authors. See Sec. 8.13.4.

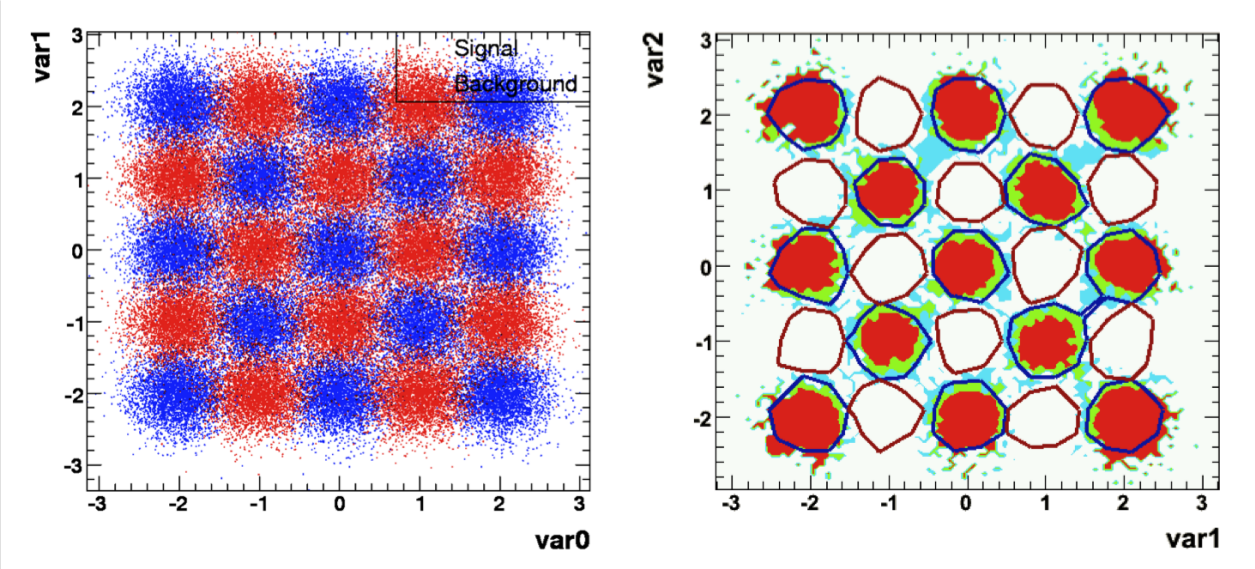

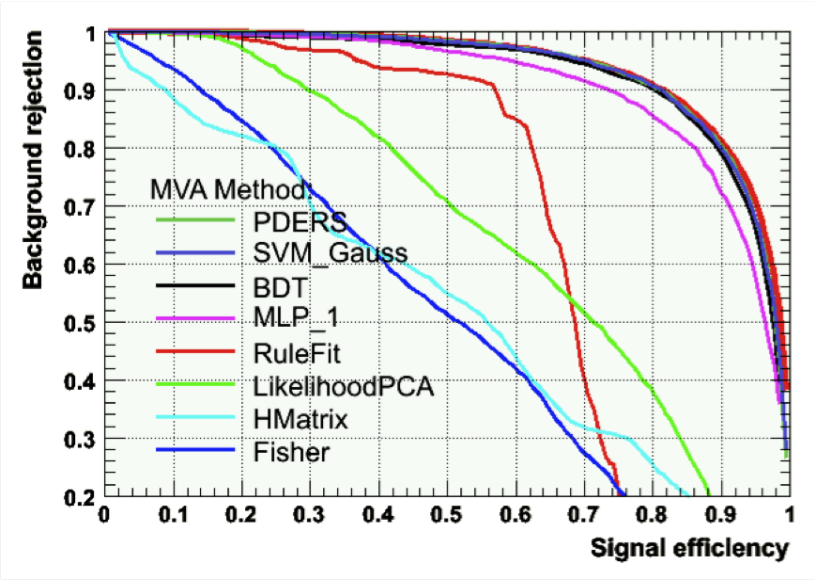

Figure 22: The top left plot shows the two-dimensional Gaussian-distributed signal (blue) and background (red) distributions for the "Schachbrett" toy Monte Carlo sample. In the top right plot is drawn as an example the Support Vector Machine's classification response in terms of the weights given to events in the two dimensional plane. Dark red signifies large MVA output values, i.e. signal like event regions. The bottom plot gives the ROC curves obtained with several TMVA classifiers for this example after training. The theoretically best ROC curve is known and shown by the thick red outer line. While all the non-linear classification methods (after carefully tuning their parameters) are able to approximate the theoretical limit, the linear methods strongly underperform.

far (see Table 7).

Individual event-weight support is now commonly realised, only missing (and not foreseen to be provided) for the less recommended neural network CFMlpANN. Support of negative event weights occurring, e.g., in NLO MC requires more scrutiny as discussed in Sec. 3.1.3 on page 19.

Ranking of the input variables cannot be defined in a straightforward manner for all MVA methods. Transparent and objective variable ranking through performance comparison of the MVA method under successive elimination of one input variable at a time is under consideration (so far only realised for the naive-Bayes likelihood classifier).

Standalone C++ response classes (not required when using the Reader application) are generated by the majority of the classifiers, but not yet for regression analysis. The missing ones for PDE-RS, PDE-Foam, k-NN, Cuts and CFMlpANN will only be considered on explicit request.

The availability of help messages, which assist the user with the performance tuning and which are printed on standard output when using the booking option 'H', is complete.

Finally, custom macros are provided for some MVA methods to analyse specific properties, such as the fidelity of likelihood reference distributions or the neural network architecture, etc. More macros can be added upon user request.

12 Conclusions and Plans

TMVA is a toolkit that unifies highly customisable multivariate (MVA) classification and regression algorithms in a single framework thus ensuring convenient use and an objective performance assessment. It is designed for machine learning applications in high-energy physics, but not restricted to these. Source code and library of TMVA-v.3.5.0 and higher versions are part of the standard ROOT distribution kit (v5.14 and higher). The newest TMVA development version can be downloaded from Sourceforge.net at http://tmva.sourceforge.net.

This Users Guide introduced the main steps of a TMVA analysis allowing a user to optimise and perform her/his own multivariate classification or regression. Let us recall the main features of the TMVA design and purpose:

- TMVA works in transparent factory mode to allow an unbiased performance assessment and comparison: all MVA methods see the same training and test data, and are evaluated following the same prescription.
- A complete TMVA analysis consists of two steps:
  1. **Training:** the ensemble of available and optimally customised MVA methods are trained and tested on independent signal and background data samples; the methods are evaluated and the most appropriate (performing and concise) ones are selected.
  2. **Application:** selected trained MVA methods are used for the classification of data samples with unknown signal and background composition, or for the estimate of unknown

MVA method	Classi-fication	Regress-ion	Multi-class/target classification	regression	Treats event weights: positive	negative	Variable ranking	Standalone response class	Help messages	Custom macros
Cut optimisation	●	○	○	○	●	○	○	○	●	○
Likelihood	●	○	○	○	●	●	●	●	●	●
PDE-RS	●	●	○	○	●	●	○	○	●	○
PDE-Foam	●	●	○	●	●	●	○	○	●	●
k-NN	●	●	○	●	●	●	○	○	●	○
H-Matrix	●	○	○	○	●	○	●	●	●	○
Fisher	●	○	○	○	●	○	●	●	●	○
LD	●	●	○	○	●	○	●	●	●	○
FDA	●	●	○	○	●	○	○	●	●	○
MLP	●	●	○	●	●	○	●	●	●	●
TMlpANN$^{(\star)}$	●	○	○	○	●	○	○	●	●	○
CFMlpANN	●	○	○	○	○	○	○	○	○	○
SVM	●	○	○	○	●	●	○	●	●	○
BDT	●	●	○	○	●	●	●	●	●	●
RuleFit	●	○	○	○	●	○	●	●	●	●

$^{(\star)}$Not a generic TMVA method − interface to ROOT class `TMultiLayerPerceptron`.

Table 7: Status of the methods with respect to various TMVA features. See text for comments. Note that the column "Standalone response class" only refers to classification. It is yet unavailable for regression.

target values (regression).

- A Factory class object created by the user organises the customisation and interaction with the MVA methods for the training, testing and evaluation phases of the TMVA analysis. The training results together with the configuration of the methods are written to result ("weight") files in XML format.

- Standardised outputs during the Factory running, and dedicated ROOT macros allow a refined assessment of each method's behaviour and performance for classification and regression.

- Once appropriate methods have been chosen by the user, they can be applied to data samples with unknown classification or target values. Here, the interaction with the methods occurs through a Reader class object created by the user. A method is booked by giving the path to its weight file resulting from the training stage. Then, inside the user's event loop, the MVA response is returned by the Reader for each of the booked MVA method, as a function of the event values of the discriminating variables used as input for the classifiers. Alternatively, for classification, the user may request from the Reader the probability that a given event belongs to the signal hypothesis and/or the event's Rarity.

- In parallel to the XML files, TMVA generates standalone C++ classes after the training, which can be used for classification problems (feature not available yet for regression). Such classes are available for all classifiers except for cut optimisation, PDE-RS, PDE-Foam, k-NN and the old CFMlpANN.

We give below a summary of the TMVA methods, outlining the current state of their implementation, their advantages and shortcomings.

- *Rectangular Cut Optimisation*
  The current implementation is mature. It includes speed-optimised range searches using binary trees, and three optimisation algorithms: Monte Carlo sampling, a Genetic Algorithm and Simulated Annealing. In spite of these tools, optimising the cuts for a large number of discriminating variables remains challenging. The user is advised to reduce the available dimensions to the most significant variables (e.g., using a principal component analysis) prior to optimising the cuts.

- *Likelihood*
  Automatic non-parametric probability density function (PDF) estimation through histogram smoothing and interpolation with various spline functions and quasi-unbinned kernel density estimators is implemented. The PDF description can be individually tuned for each input variable.

- *PDE-RS*
  The multidimensional probability density estimator (PDE) approach is in an advanced development stage featuring adaptive range search, several kernel estimation methods, and speed optimised range search using event sorting in binary trees. It has also been extended to regression.

- *PDE-Foam*
  This new multidimensional PDE algorithm uses self-adapting phase-space binning and is a fast realisation of PDE-RS in fixed volumes, which are determined and optimised during the training phase. Much work went into the development of PDE-Foam. It has been thoroughly tested, and can be considered a mature method. PDE-Foam performs classification and regression analyses.

- *k-NN*
  The k-Nearest Neighbour classifier is also in a mature state, featuring both classification and regression. The code has been well tested and shows satisfactory results. With scarce training statistics it may slightly underperform in comparison with PDE-RS, whereas it is significantly faster in the application to large data samples.

- *Fisher and H-Matrix*
  Both are mature algorithms, featuring linear discrimination for classification only. Higher-order correlations are taken care of by FDA (see below).

- *Linear Discriminant (LD)*
  LD is equivalent to Fisher but providing both classification and linear regression.

- *Function Discriminant Analysis (FDA)*
  FDA is a mature algorithm, which has not been extensively used yet. It extends the linear discriminant to moderately non-linear correlations that are fit to the training data.

- *Artificial Neural Networks*
  Significant work went into the implementation of fast feed-forward multilayer perceptron algorithms into TMVA. Two external ANNs have been integrated as fully independent methods, and another one has been newly developed for TMVA, with emphasis on flexibility and speed. The performance of the latter ANN (MLP) has been cross checked against the Stuttgart ANN (using as an example τ identification in ATLAS), and was found to achieve competitive performance. The MLP ANN also performs multi-target regression.

- *Support Vector Machine*
  SVM is a relatively new multivariate analysis algorithm with a strong statistical background. It performs well for nonlinear discrimination and is insensitive to overtraining. Optimisation is straightforward due to a low number of adjustable parameters (only two in the case of Gaussian kernel). The response speed is slower than for a not-too-exhaustive neural network, but comparable with other nonlinear methods. SVM is being extended to multivariate regression.

- *Boosted Decision Trees*
  The BDT implementation has received constant attention over the years of its development. The current version includes additional features like bagging or gradient boosting, and manual or automatic pruning of statistically insignificant nodes. It is a highly performing MVA method that also applies to regression problems.

- *RuleFit*
  The current version has the possibility to run either the original program written by J. Friedman [32] or an independent TMVA implementation. The TMVA version has been improved

both in speed and performance and achieves almost equivalent results with respect to the original one, requiring however somewhat more tuning.

The new framework introduced with TMVA 4 provides the flexibility to combine MVA methods in a general fashion. Exploiting these capabilities for classification and regression however requires to create so-called committee methods for each combination. So far, we provide a generalised Boost method, allowing to boost any classifier by simply setting the variable `Boost_Num` in the configuration options to a positive number (plus possible adjustment of other configuration parameters). The result is a potentially powerful committee method unifying the excellent properties of boosting with MVA methods that already represent highly optimised algorithms.

Boosting is not the only combination the new framework allows us to establish. TMVA now also supports the categorisation of classification according to the phase space, which allows a better modelling and hence simplifies the mapping of the feature space. This particularly improves simple methods, such las projective likelihood and linear discriminants.

Acknowledgements

The fast growth of TMVA would not have been possible without the contribution and feedback from many developers (also co-authors of this Users Guide) and users to whom we are indebted. We thank in particular the CERN Summer students Matt Jachowski (Stanford U.) for the implementation of TMVA's MLP neural network, and Yair Mahalalel (Tel Aviv U.) for a significant improvement of PDE-RS. The Support Vector Machine has been contributed to TMVA by Andrzej Zemla and Marcin Wolter (IFJ PAN Krakow), and the k-NN method has been written by Rustem Ospanov (Texas U.). We are grateful to Lucian Ancu, Doug Applegate, Kregg Arms, René Brun and the ROOT team, Andrea Bulgarelli, Marc Escalier, Zhiyi Liu, Colin Mclean, Elzbieta Richter-Was, Alfio Rizzo, Lydia Roos, Vincent Tisserand, Alexei Volk, Jiahang Zhong for helpful feedback and bug reports. Thanks also to Lucian Ancu for improving the plotting macros.

A More Classifier Booking Examples

The Code Examples 53 and 54 give a (non-exhaustive) collection of classifier bookings with appropriate default options. They correspond to the example training job `TMVAClassification.C`.

```
// Cut optimisation using Monte Carlo sampling
factory->BookMethod( TMVA::Types::kCuts, "Cuts",
    "!H:!V:FitMethod=MC:EffSel:SampleSize=200000:VarProp=FSmart" );

// Cut optmisation using Genetic Algorithm
factory->BookMethod( TMVA::Types::kCuts, "CutsGA",
    "H:!V:FitMethod=GA:CutRangeMin=-10:CutRangeMax=10:VarProp[1]=FMax:EffSel:\
     Steps=30:Cycles=3:PopSize=400:SC_steps=10:SC_rate=5:SC_factor=0.95" );

// Cut optmisation using Simulated Annealing algorithm
factory->BookMethod( TMVA::Types::kCuts, "CutsSA",
    "!H:!V:FitMethod=SA:EffSel:MaxCalls=150000:KernelTemp=IncAdaptive:\
     InitialTemp=1e+6:MinTemp=1e-6:Eps=1e-10:UseDefaultScale" );

// Likelihood classification (naive Bayes) with Spline PDF parametrisation
factory->BookMethod( TMVA::Types::kLikelihood, "Likelihood",
    "H:!V:TransformOutput:PDFInterpol=Spline2:NSmoothSig[0]=20:\
     NSmoothBkg[0]=20:NSmoothBkg[1]=10:NSmooth=1:NAvEvtPerBin=50" );

// Likelihood with decorrelation of input variables
factory->BookMethod( TMVA::Types::kLikelihood, "LikelihoodD",
    "!H:!V:!TransformOutput:PDFInterpol=Spline2:NSmoothSig[0]=20:\
     NSmoothBkg[0]=20:NSmooth=5:NAvEvtPerBin=50:VarTransform=Decorrelate" );

// Likelihood with unbinned kernel estimator for PDF parametrisation
factory->BookMethod( TMVA::Types::kLikelihood, "LikelihoodKDE",
    "!H:!V:!TransformOutput:PDFInterpol=KDE:KDEtype=Gauss:KDEiter=Adaptive:\
     KDEFineFactor=0.3:KDEborder=None:NAvEvtPerBin=50" );
```

Code Example 53: Examples for booking MVA methods in TMVA for application to classification and – where available – to regression problems. The first argument is a unique type enumerator (the avaliable types can be looked up in `src/Types.h`), the second is a user-defined name (must be unique among all booked classifiers), and the third a configuration option string that is specific to the classifier. For options that are not set in the string default values are used. The syntax of the options should become clear from the above examples. Individual options are separated by a ':'. Boolean variables can be set either explicitly as `MyBoolVar=True/False`, or just via `MyBoolVar/!MyBoolVar`. All concrete option variables are explained in the tools and classifier sections of this Users Guide. The list is continued in Code Example 54.

```
// Probability density estimator range search method (multi-dimensional)
factory->BookMethod( TMVA::Types::kPDERS, "PDERS",
    "!H:V:NormTree=T:VolumeRangeMode=Adaptive:KernelEstimator=Gauss:\
     GaussSigma=0.3:NEventsMin=400:NEventsMax=600" );

// Multi-dimensional PDE using self-adapting phase-space binning
factory->BookMethod( TMVA::Types::kPDEFoam, "PDEFoam",
    "H:V:SigBgSeparate=F:TailCut=0.001:VolFrac=0.0333:nActiveCells=500:\
     nSampl=2000:nBin=5:CutNmin=T:Nmin=100:Kernel=None:Compress=T" );

// k-Nearest Neighbour method (similar to PDE-RS)
factory->BookMethod( TMVA::Types::kKNN, "KNN",
    "H:nkNN=20:ScaleFrac=0.8:SigmaFact=1.0:Kernel=Gaus:UseKernel=F:\
     UseWeight=T:!Trim" );

// H-matrix (chi-squared) method
factory->BookMethod( TMVA::Types::kHMatrix, "HMatrix", "!H:!V" );

// Fisher discriminant (also creating Rarity distribution of MVA output)
factory->BookMethod( TMVA::Types::kFisher, "Fisher",
    "H:!V:Fisher:CreateMVAPdfs:PDFInterpolMVAPdf=Spline2:NbinsMVAPdf=60:\
     NsmoothMVAPdf=10" );

// Fisher discriminant with Gauss-transformed input variables
factory->BookMethod( TMVA::Types::kFisher, "FisherG", "VarTransform=Gauss" );

// Fisher discriminant with principle-value-transformed input variables
factory->BookMethod( TMVA::Types::kFisher, "FisherG", "VarTransform=PCA" );

// Boosted Fisher discriminant
factory->BookMethod( TMVA::Types::kFisher, "BoostedFisher",
    "Boost_Num=20:Boost_Transform=log:\
     Boost_Type=AdaBoost:Boost_AdaBoostBeta=0.2");

// Linear discriminant (same as Fisher, but also performing regression)
factory->BookMethod( TMVA::Types::kLD, "LD", "H:!V:VarTransform=None" );

// Function discrimination analysis (FDA), fitting user-defined function
factory->BookMethod( TMVA::Types::kFDA, "FDA_MT",
    "H:!V:Formula=(0)+(1)*x0+(2)*x1+(3)*x2+(4)*x3:\
     ParRanges=(-1,1);(-10,10);(-10,10);(-10,10);(-10,10):FitMethod=MINUIT:\
     ErrorLevel=1:PrintLevel=-1:FitStrategy=2:UseImprove:UseMinos:SetBatch" );
```

Code Example 54: Continuation from Code Example 53. Continued in Code Example 54.

```
// Artificial Neural Network (Multilayer perceptron) - TMVA version
factory->BookMethod( TMVA::Types::kMLP, "MLP",
    "H:!V:NeuronType=tanh:VarTransform=N:NCycles=600:HiddenLayers=N+5:\
     TestRate=5" );

// NN with BFGS quadratic minimisation
factory->BookMethod( TMVA::Types::kMLP, "MLPBFGS",
    "H:!V:NeuronType=tanh:VarTransform=N:NCycles=600:HiddenLayers=N+5:\
     TestRate=5:TrainingMethod=BFGS" );

// NN (Multilayer perceptron) - ROOT version
factory->BookMethod( TMVA::Types::kTMlpANN, "TMlpANN",
    "!H:!V:NCycles=200:HiddenLayers=N+1,N:LearningMethod=BFGS:
     ValidationFraction=0.3"  );

// NN (Multilayer perceptron) - ALEPH version (depreciated)
factory->BookMethod( TMVA::Types::kCFMlpANN, "CFMlpANN",
    "!H:!V:NCycles=2000:HiddenLayers=N+1,N"  );

// Support Vector Machine
factory->BookMethod( TMVA::Types::kSVM, "SVM", "Gamma=0.25:Tol=0.001" );

// Boosted Decision Trees with adaptive boosting
factory->BookMethod( TMVA::Types::kBDT, "BDT",
    "!H:!V:NTrees=400:nEventsMin=400:MaxDepth=3:BoostType=AdaBoost:\
     SeparationType=GiniIndex:nCuts=20:PruneMethod=NoPruning" );

// Boosted Decision Trees with gradient boosting
factory->BookMethod( TMVA::Types::kBDT, "BDTG",
    "!H:!V:NTrees=1000:BoostType=Grad:Shrinkage=0.30:UseBaggedGrad:\
     GradBaggingFraction=0.6:SeparationType=GiniIndex:nCuts=20:\
     PruneMethod=CostComplexity:PruneStrength=50:NNodesMax=5" );

// Boosted Decision Trees with bagging
factory->BookMethod( TMVA::Types::kBDT, "BDTB",
    "!H:!V:NTrees=400:BoostType=Bagging:SeparationType=GiniIndex:\
     nCuts=20:PruneMethod=NoPruning" );

// Predictive learning via rule ensembles (RuleFit)
factory->BookMethod( TMVA::Types::kRuleFit, "RuleFit",
    "H:!V:RuleFitModule=RFTMVA:Model=ModRuleLinear:MinImp=0.001:\
     RuleMinDist=0.001:NTrees=20:fEventsMin=0.01:fEventsMax=0.5:\
     GDTau=-1.0:GDTauPrec=0.01:GDStep=0.01:GDNSteps=10000:GDErrScale=1.02" );
```

Code Example 55: Continuation from Code Example 54.

References

[1] R. Brun and F. Rademakers, "*ROOT - An Object Oriented Data Analysis Framework*", Nucl. Inst. Meth. in Phys. Res. A 389, 81 (1997).

[2] J. Friedman, T. Hastie and R. Tibshirani, "*The Elements of Statistical Learning*", Springer Series in Statistics, 2001.

[3] A. Webb, "*Statistical Pattern Recognition*", 2nd Edition, J. Wiley & Sons Ltd, 2002.

[4] L.I. Kuncheva, "*Combining Pattern Classifiers*", J. Wiley & Sons, 2004.

[5] I. Narsky, "*StatPatternRecognition: A C++ Package for Statistical Analysis of High Energy Physics Data*", physics/0507143 (2005).

[6] The following web pages give information on available statistical tools in HEP and other areas of science: https://plone4.fnal.gov:4430/P0/phystat/, http://astrostatistics.psu.edu/statcodes/.

[7] The BABAR Physics Book, BABAR Collaboration (P.F. Harrison and H. Quinn (editors) *et al.*), SLAC-R-0504 (1998); S. Versillé, PhD Thesis at LPNHE, http://lpnhe-babar.in2p3.fr/theses/these_SophieVersille.ps.gz (1998).

[8] To our information, the *Rarity* has been originally defined by F. Le Diberder in an unpublished Mark II internal note. In a single dimension, as defined in Eq. (13), it is equivalent to the μ-transform developed in: M. Pivk, "*Etude de la violation de CP dans la désintégration* $B^0 \to h^+h^-$ $(h = \pi, K)$ *auprès du détecteur BABAR à SLAC*", PhD thesis (in French), http://tel.archives-ouvertes.fr/documents/archives0/00/00/29/91/index_fr.html (2003).

[9] M. Kendall, A. Stuart, K.J. Ord, and S. Arnold, Kendall's Advanced Theory of Statistics: Volume 2A – Classical Inference and and the Linear Model (Kendall's Library of Statistics), A Hodder Arnold Publication, 6th edition, April 1999.

[10] T.M. Cover and J.A. Thomas, "*Elements of information theory*", Wiley-Interscience, New York, NY, USA, 1991.

[11] Y.-I. Moon, B. Rajagopalan, and U. Lall, Phys. Rev. E, 52 (no. 3), 2318 (1995).

[12] J.L. Bentley, "*Multidimensional Binary Search Trees Used for Associate Searching*", Communications of the ACM, 18, 509 (1975); J.L. Bentley, "*Multidimensional Divide-and-Conquer*", Communications of the ACM, 23(4), 214 (1980); J.H. Friedman, J.L. Bentley and R.A. Finkel, "*An Algorithm for Finding Matches in Logarithmic Expected Time*", ACM Trans. on Mathematical Software, 3(3), 209 (1977).

[13] R. Sedgewick, "*Algorithms in C++*", Addison Wesley, Chapter 26, Boston, USA (1992).

[14] T. Carli and B. Koblitz, Nucl. Instrum. Meth. A 501, 576 (2003) [hep-ex/0211019].

[15] S. Jadach, Comput. Phys. Commun. 130, 244 (2000); S. Jadach, Comput. Phys. Commun. 152, 55 (2003).

[16] D. Dannheim et al., Nucl. Instr. and Meth. A 606, 717 (2009).

[17] F. James, "*MINUIT, Function Minization and ERROR Analysis*", Version 94.1, CERN program Library long writeup D506.

[18] P.J.M. Van Laarhoven and E.H.L. Aart, "*Simulated Annealing: Theory and Application*", D. Reidel Publishing, Dordrecht, Holland, 1987.

[19] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.M. Teller, and E. Teller, J. Chem. Phys. 21, 6, 1087 (1953).

[20] 353QH twice smoothing algorithm, presented by J. Friedman in Proc. of the 1974 CERN School of Computing, Norway, Aug 11-24, 1974.

[21] D.W. Scott, "*Multivariate Density Estimation, Theory, Practice, and Visualization*", Wiley-Interscience, New York, 1992.

[22] R.A. Fisher, Annals Eugenics 7, 179 (1936).

[23] P.C. Mahalanobis, Proc. Nat. Inst. Sci. India, Part 2A, 49 (1936); P.C. Mahalanobis, "*On the generalized distance in statistics*", Proceedings of the National Institute of Science, Calcutta, 12, 49 (1936).

[24] C.G. Broyden, "*The Convergence of a Class of Double-rank Minimization Algorithms*", J. Inst. of Math. and App. 6, 76 (1970); R. Fletcher, "*A New Approach to Variable Metric Algorithms*", Computer J. 13, 317 (1970); D. Goldfarb, "*A Family of Variable Metric Updates Derived by Variational Means*", Math. Comp. 24, 23 (1970); D.F. Shannon, "*Conditioning of Quasi-Newton Methods for Function Minimization*", Math. Comp. 24, 647 (1970).

[25] R.E. Schapire, "*The strenght of weak learnability*, (1990) Machine Learning 5 197-227; Y. Freud, "*Boosting a weak learning algorithm by majority* (1995) Inform. and Comput. **121** 256-285;

[26] Y. Freund and R.E. Schapire, J. of Computer and System Science 55, 119 (1997).

[27] H. Drucker,*Improving regressors using boosting techniques*,In D. H. Fisher (Ed.), Proceedings of the fourteenth international conference on machine learning (ICML 1997) (pp. 107115). Nashville, TN, USA, July 812. Morgan Kaufmann, ISBN 1558604863.

[28] L. Breiman, *Random Forests*, Technical Report, University of California 2001.

[29] Y.R. Quinlan, "*Simplifying Decision Trees*", Int. J. Man-Machine Studies, 28, 221 (1987).

[30] L. Breiman, J. Friedman, R. Olshen and C. Stone, "*Classification and Regression Trees*", Wadsworth (1984).

[31] We use the ROOT class `TParallelCoord` by B. Dalla Piazza, 2007.

[32] J. Friedman and B.E. Popescu, "*Predictive Learning via Rule Ensembles*", Technical Report, Statistics Department, Stanford University, 2004.

[33] J. Friedman and B.E. Popescu, "*Gradient Directed Regularization for Linear Regression and Classification*", Technical Report, Statistics Department, Stanford University, 2003.

[34] RuleFit web site: http:/www-stat.stanford.edu/~jhf/r-rulefit/RuleFit_help.html.

[35] J. Conrad and F. Tegenfeldt, JHEP 0607, 040 (2006) [hep-ph/0605106].

[36] W. Cohen and Y. Singer, Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99), 335, AAAI Press, 1999.

[37] V. Vapnik and A. Lerner, "*Pattern recognition using generalized portrait method*", Automation and Remote Control, 24, 774 (1963).

[38] V. Vapnik and A. Chervonenkis, "*A note on one class of perceptrons*", Automation and Remote Control, 25 (1964).

[39] B.E. Boser, I.M. Guyon, and V.N. Vapnik, "*A training algorithm for optimal margin classifiers*", in D. Haussler, ed., Proceedings of the 5th Annual ACM Workshop on Computational Leraning Theory, 144, ACM Press (1992).

[40] C. Cortes and V. Vapnik, "*Support vector networks*", Machine Learning, 20, 273 (1995).

[41] V. Vapnik, "*The Nature of Statistical Learning Theory*", Springer Verlag, New York, 1995.

[42] C.J.C. Burges, "*A Tutorial on Support Vector Machines for Pattern Recognition*", Data Mining and Knowledge Discovery, 2, 1 (1998).

[43] J. Platt, "*Fast training of support vector machines using sequential minimal optimization*", in B. Scholkopf, C. Burges and A. Smola, eds., Advances in Kernel Methods – Support Vector Learning, ch. 12, pp. 185, MIT Press, 1999.

[44] S. Keerthi, S. Shevade, C. Bhattacharyya and K. Murthy, "*Improvements to Platt's SMO algorithm for SVM classifier design*", Technical Report CD-99-14, Dept. of Mechanical and Production Engineering, Natl. Univ. Singapore, Singapore, 1999.

[45] P. Huber, "*Robust estimation of a location parameter*", Annals of Mathematical Statistics, 35, 73-101,1964

Index

AdaBoost, 52
AdaBoost.R2, 53
Adaptive boost, 52
Artificial neural networks, 91

Back propagation, 97
Bagging, 55
BFGS, 98
Binary search trees, 41, 59
Binary split, 57
Booking variable transformations, 40
Boosted, 119
Boosting, 52, 119
    Adaptive Boost, 52
    Gradient Boost, 54
Boosting method, 118

Category method, 118
Chi-squared estimator, 83
Classification probability, 28
Clermont-Ferrand neural network, 91
Combining, 118
Committee method, 118
Copyright, 2
Correlation, 24
Correlation ratio, 25
Cross Entropy, 109
Cut optimisation, 41, 57
Cuts, 57

Discriminating variables
    decorrelation of, 38
    Gaussianisation of, 40
    normalisation of, 38
    PCA of, 39
    preprocessing of, 37
    transformation of, 37
Download
    from Sourceforge.net, 4
    from SVN, 4

Evaluating MVA methods, 23
Examples, 5

Factory, 5, 13, 14
    booking MVA methods, 21
    preparing training and test data, 19
    selecting input variables, 17
    specifying event weights, 19
    specifying input data (trees), 15
FDA, 88
Feed-forward MLP, 91
Fisher discriminant, 84
Fitness, 47
    for cut optimisation, 61
Fitting, 45
Forest, 104, 107
Function Discriminant Analysis, 88

Genetic Algorithm, 47, 61
    combination with Minuit, 51
Gini Index, 109
Graphical user interface (GUI), 7, 29

H-Matrix, 83
Help
    booking options, 22
    method-specific help messages, 22
    MVA method optimisation, 22
Histogram smoothing, 42
Huber loss function, 55

k-Nearest Neighbour classifier, 79
k-NN, 79
kd-tree, 41
Kernel density estimators (KDE), 44
Kernel estimation, 67, 73

LD, 86
License, 3
Likelihood, 62
Likelihood output transformation, 63
Linear Discriminant, 86
Linear Discriminant Analysis, 84

Mahalanobis distance, 85
Makefile, 4
Minuit

combination with MC or GA, 51
cut optimisation, 59
fitter, 46
minimisation, 46
Misclassification error, 109
MLP neural network, 93
Monte Carlo sampling, 46, 60
combination with Minuit, 51
Multilayer perceptron (MLP), 91

Negative Event Weights, 19
Nonlinear discriminant analysis, 91
Nonseparable data, 102

Overtraining, 27

PDE, 41, 64
PDE-Foam, multi-dimensional likelihood, 69
PDE-RS, 41
PDE-RS, multidimensional likelihood, 64
PDF, 41, 62
PDF parameterisation
with kernel density estimators, 44
with splines, 42
Performance evaluation
background rejection vs. signal efficiency, 24
Performance evaluation
significance, 24
Performance evaluation, 23
background rejection vs. signal efficiency, 12
separation, 24
Principal component decomposition, analysis, 39
Probability Density Function, 41, 62
Pruning, 109

Randomising, 108
Rarity, 28
Reader, 5, 14, 31
booking MVA methods, 33
requesting the MVA response, 33
requesting the Rarity of a classifier, 34
requesting the signal probability of a classifier, 34
specifying input variables, 31
Receiver Operating Characteristic (ROC), 7
Rectangular cut optimisation, 57
Rectangular cuts, 57
Regression
multi-target (multidimensional), 18
Resampling, 55
ROOT
compatibility, 5
macros, 29–31
ROOT neural network, 92
RuleFit, 111
Rules
base learners, 112
distance between, 114
ensemble of, 111
fit path, 115
fitting of, 115
generation of, 114
linear terms, 112
loss function of, 115
risk of, 115
variable importance, 116

Signal probability, 28
Simulated Annealing, 50, 61
Spectator variables, 18
Splines, 42, 43
Standalone C++ response classes, 34
Support vector machine, SVM, 99
linear, 101
nonlinear, 102
performance of, 104

Testing multivariate methods, 23
TMultiLayerPerceptron, 92
TMVA analysis flow, 14
TMVAClassification, 5, 12
TMVAClassificationApplication, 5, 12
TMVAGui, 8
TMVARegression, 5, 12
TMVARegressionApplication, 5, 12
Training MVA methods, 22

Weight files, 22
Weight files

XML format, 22, 33