

# Machine learning techniques available in pRoloc

LAURENT GATTO\*

Cambridge Center for Proteomics  
University of Cambridge

February 4, 2014

## Contents

|          |   |          |
|----------|---|----------|
| <b>1</b> | <b>Introduction</b>   | <b>2</b> |
| <b>2</b> | <b>Data sets</b>  | <b>2</b> |
| <b>3</b> | <b>Unsupervised machine learning</b>                          | <b>3</b> |
| <b>4</b> | <b>Supervised machine learning</b>                            | <b>3</b> |
| 4.1      | Algorithms used . . . . .                                     | 4        |
| 4.2      | Default analysis scheme . . . . .                             | 5        |
| 4.3      | Customising model parameters . . . . .                        | 7        |
| <b>5</b> | <b>Comparison of different classifiers</b>                    | <b>7</b> |
| <b>6</b> | <b>Semi-supervised machine learning and novelty detection</b> | <b>8</b> |

---

\*lg390@cam.ac.uk

**NB** This document is currently under construction; please see the `pRoloc` tutorial for details.

**TODO** Define nomenclature

## 1 Introduction

For a general practical introduction to `pRoloc`, readers are referred to the tutorial, available using `vignette("pRoloc-tutorial", package = "pRoloc")`. The following document provides a overview of the algorithms available in the package. The respective section describe unsupervised machine learning (USML), supervised machine learning (SML), and semi-supervised machine learning (SSML) as implemented in the novelty detection algorithm.

## 2 Data sets

We provide 12 test data sets in the `pRolocdata` package that can be readily used with `pRoloc`. The data set can be listed with `pRolocdata` and loaded with the `data` function. Each data set, including its origin, is individually documented.

The data sets are distributed as `MSnSet` instances. Briefly, these are dedicated containers for quantitation data as well as feature and sample meta-data. More details about `MSnSets` are available in the `pRoloc` tutorial and in the `MSnbase` package, that defined the class.

```
> library("pRolocdata")
> data(tan2009r1)
> tan2009r1

MSnSet (storageMode: lockedEnvironment)
assayData: 888 features, 4 samples
  element names: exprs
protocolData: none
phenoData
  sampleNames: X114 X115 X116 X117
  varLabels: Fractions
  varMetadata: labelDescription
featureData
  featureNames: CG10060 CG10067 ... CG9983 (888
  total)
```

```
fvarLabels: FBgn Flybase.Symbol ... pd.markers
(9 total)
fvarMetadata: labelDescription
experimentData: use 'experimentData(object) '
pubMedIds: 19317464
Annotation:
- - - Processing information - - -
MSnbase version: 1.7.24
```

### 3 Unsupervised machine learning

Unsupervised machine learning refers to clustering, i.e. finding structure in a quantitative, generally multi-dimensional data set of unlabelled data.

Currently, unsupervised clustering facilities are available through the `plot2D` function and the `MLInterfaces` package [Carey et al.](#). The former takes an `MSnSet` instance and represents the data on a scatter plot along the first two principal components. Arbitrary feature meta-data can be represented using different colours and point characters. The reader is referred to the manual page available through `?plot2D` for more details and examples.

`pRoloc` also implements a `MLean` method for `MSnSet` instances, allowing to use the relevant infrastructure with the organelle proteomics framework. Although provides a common interface to unsupervised and numerous supervised algorithms, we refer to the `pRoloc` tutorial for its usage to several clustering algorithms.

### 4 Supervised machine learning

Supervised machine learning refers to a broad family of classification algorithms. The algorithms learns from a modest set of labelled data points called the training data. Each training data example consists of a pair of inputs: the actual data, generally represented as a vector of numbers and a class label, representing the membership to exactly 1 of multiple possible classes. When there are only two possible classes, one refers to binary classification. The training data is used to construct a model that can be used to classifier new, unlabelled examples. The model takes the numeric vectors of the unlabelled data points and return, for each of these inputs, the corresponding mapped class.

## 4.1 Algorithms used

**k-nearest neighbour** Function `knn` from package `class`. For each row of the test set, the  $k$  nearest (in Euclidean distance) training set vectors are found, and the classification is decided by majority vote over the  $k$  classes, with ties broken at random. This is a simple algorithm that is often used as baseline classifier.

**Partial least square DA** Function `plsda` from package `caret`. Partial least square discriminant analysis is used to fit a standard PLS model for classification.

**Support vector machine** A support vector machine constructs a hyperplane (or set of hyperplanes for multiple-class problem), which are then used for classification. The best separation is defined as the hyperplane that has the largest distance (the margin) to the nearest data points in any class, which also reduces the classification generalisation error. To assure linear separation of the classes, the data is transformed using a *kernel function* into a high-dimensional space, permitting linear separation of the classes.

`pRoloc` makes use of the functions `svm` from package `e1071` and `ksvm` from `kernelab`.

**Artificial neural network** Function `nnet` from package `nnet`. Fits a single-hidden-layer neural network, possibly with skip-layer connections.

**Naive Bayes** Function `naiveBayes` from package `e1071`. Naive Bayes classifier that computes the conditional a-posterior probabilities of a categorical class variable given independent predictor variables using the Bayes rule. Assumes independence of the predictor variables, and Gaussian distribution (given the target class) of metric predictors.

**Random Forest** Function `randomRandom` from package `randomForest`.

**Chi-square** Assignment based on squared differences between a labelled marker and a new feature to be classified. Canonical protein correlation profile method (PCP) uses squared differences between a labelled marker and new features. In (Andersen et al., 2003),  $\chi^2$  is defined as *the [summed] squared deviation of the normalized profile [from the marker] for all peptides divided by the number of data points*, i.e.  $\chi^2 = \frac{\sum_{i=1}^n (x_i - m_i)^2}{n}$ , whereas (Wiese et al., 2007) divide by the value the squared value by the value of the reference feature in each fraction, i.e.  $\chi^2 = \sum_{i=1}^n \frac{(x_i - m_i)^2}{m_i}$ , where  $x_i$  is normalised intensity of feature  $x$  in fraction  $i$ ,  $m_i$  is the normalised intensity of marker  $m$  in fraction  $i$  and  $n$  is the number of fractions available. We will use the former definition.

**PerTurbo** From [Courty et al. \(2011\)](#): PerTurbo, an original, non-parametric and efficient classification method is presented here. In our framework, the manifold of each class is characterized by its Laplace-Beltrami operator, which is evaluated with classical methods involving the graph Laplacian. The classification criterion is established thanks to a measure of the magnitude of the spectrum perturbation of this operator. The first experiments show good performances against classical algorithms of the state-of-the-art. Moreover, from this measure is derived an efficient policy to design sampling queries in a context of active learning. Performances collected over toy examples and real world datasets assess the qualities of this strategy.

The PerTurbo implementation comes from the `pRoloc` packages.

## 4.2 Default analysis scheme

We present below a typical classification using `pRoloc`. The analysis typically consists of two steps. The first one is to optimise the classifier parameters to be used for training and testing (see beginning of this section). A range of parameters are tested using the labelled data, for which the labels are known. For each set of parameters, we hide the labels of a subset of labelled data and use the other part to train a model and apply in on the data with hidden labels. The comparison of the estimated and expected labels enables to assess the validity of the model and hence the adequacy if the parameters. Once adequate parameters have been identified, they are used to infer a model on the complete test set and use of to infer the labels of the unlabelled examples.

### Parameter optimisation

Algorithmic performance is estimated using a stratified 20/80 partitioning. The 80% partitions are subjected to 5-fold cross-validation in order to optimise free parameters via a grid search, and these parameters are then applied to the remaining 20%. The procedure is repeated  $n = 100$  times to sample  $n$  accuracy metrics (see below) values using  $n$ , possibly different, optimised parameters for evaluation.

Models accuracy is evaluated using the F1 score,  $F1 = 2 \frac{precision \times recall}{precision + recall}$ , calculated as the harmonic mean of the precision ( $precision = \frac{tp}{tp+fp}$ , a measure of *exactness* – returned output is a relevant result) and recall ( $recall = \frac{tp}{tp+fn}$ , a measure of *completeness* – indicating how much was missed from the output). What we are aiming for are high generalisation accuracies, i.e high  $F1$ , indicating that the marker proteins in the test data set are consistently correctly assigned by the algorithms.

The result of the optimisation procedure are stored in an `GenRegRes` object that can be inspected, plotted and best parameter pairs can be extracted.

For a given algorithm `alg`, the corresponding parameter optimisation function is names `algOptimisation` or, equivalently, `olgOptimization`. See [table 1](#) for

details. A description of each of the respective model parameters is provided in the optimisation function manuals, available through `?algOptimisation`.

```
> params <- svmOptimisation(tan2009r1, times = 10, xval = 5,
+   verbose = FALSE)
> params

Object of class "GenRegRes"
Algorithm: svm
Hyper-parameters:
  cost: 0.0625 0.125 0.25 0.5 1 2 4 8 16
  sigma: 0.01 0.1 1 10 100 1000
Design:
  Replication: 10 x 5-fold X-validation
  Partitioning: 0.2/0.8 (test/train)
Results
macro F1:
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
  0.889  0.889  1.000  0.956  1.000  1.000
best sigma: 1 0.1
best cost: 1 0.5
```

## Classification

```
> tan2009r1 <- svmClassification(tan2009r1, params)
> tan2009r1

MSnSet (storageMode: lockedEnvironment)
assayData: 888 features, 4 samples
  element names: exprs
protocolData: none
phenoData
  sampleNames: X114 X115 X116 X117
  varLabels: Fractions
  varMetadata: labelDescription
featureData
  featureNames: CG10060 CG10067 ... CG9983 (888
  total)
  fvarLabels: FBgn Flybase.Symbol ... svm.scores
  (11 total)
```

```

fvarMetadata: labelDescription
experimentData: use 'experimentData(object)'
pubMedIds: 19317464
Annotation:
- - - Processing information - - -
Performed svm prediction (sigma=1 cost=1) Tue Feb  4 22:20:10 2014
MSnbase version: 1.7.24

```

### 4.3 Customising model parameters

Below we illustrate how to weight different classes according to the number of labelled instances, where large sets are down weighted. This strategy can help with imbalanced designs.

```

> w <- table(fData(dunkley2006)$markers)
> w <- 1/w[-5]
> wpar <- svmOptimisation(dunkley2006, class.weights = w)
> wres <- svmClassification(dunkley2006, pw, class.weights = w)

```


| parameter optimisation | classification         | algorithm              | package      |
|------------------------|------------------------|------------------------|--------------|
| knnOptimisation        | knnClassification      | nearest neighbour      | class        |
| ksvmOptimisation       | ksvmClassification     | support vector machine | kernlab      |
| nbOptimisation         | nbClassification       | naive bayes            | e1071        |
| nnetOptimisation       | nnetClassification     | neural networks        | nnet         |
| perTurboOptimisation   | perTurboClassification | PerTurbo               | pRoloc       |
| plsdaOptimisation      | plsdaClassification    | partial least square   | caret        |
| rfOptimisation         | rfClassification       | random forest          | randomForest |
| svmOptimisation        | svmClassification      | support vector machine | e1071        |

**Table 1:** Supervised ML algorithm available in pRoloc.

## 5 Comparison of different classifiers

Several supervised machine learning algorithms have already been applied to organelle proteomics data classification: partial least square discriminant analysis in [Dunkley et al. \(2006\)](#); [Tan et al. \(2009\)](#), support vector machines (SVMs) in [Trotter et al. \(2010\)](#), random forest in [Ohta et al. \(2010\)](#), neural networks in [Tardif et al. \(2012\)](#), naive Bayes [Nikolovski et al. \(2012\)](#). In our HUPO 2011 poster (see `vignette("HUPO.2011_poster", package = "pRoloc")`), we show that different classification algorithms provide very similar performance. We have extended this

comparison on various datasets distributed in the pRolocdata package. On figure 1, we illustrate how different algorithms reach very similar performances on most of our test datasets.

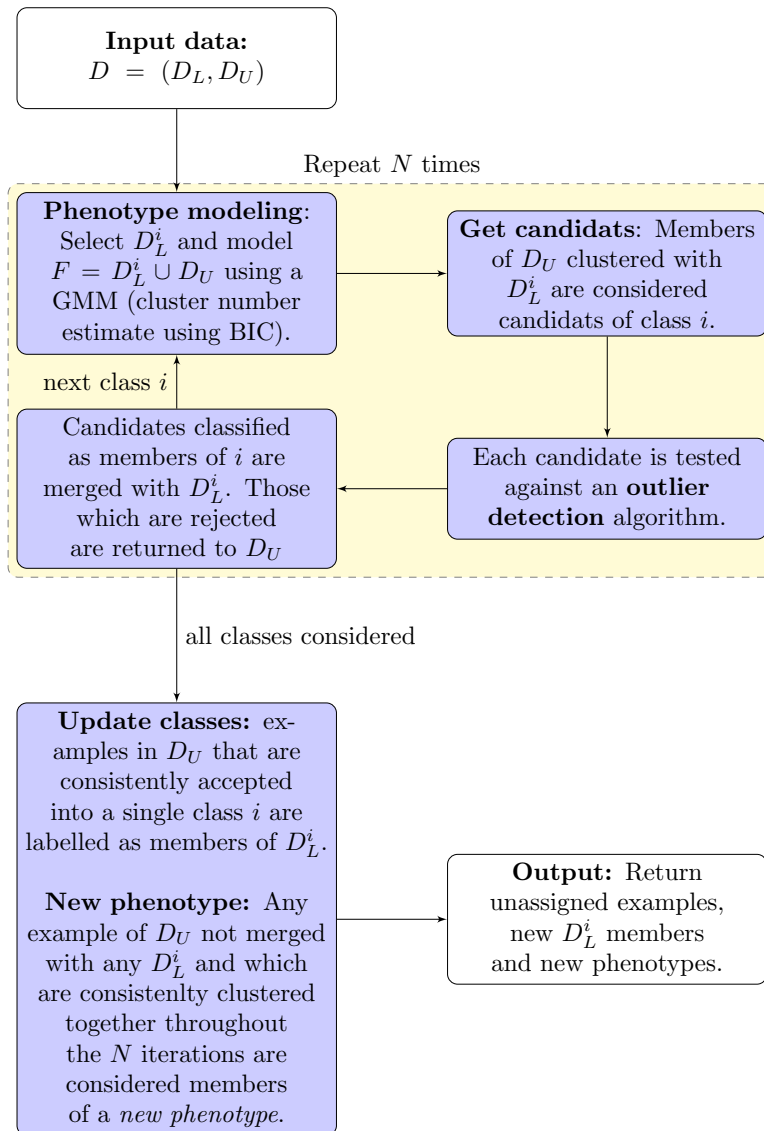


**Figure 1:** Comparison of classification performances of several contemporary classification algorithms on data from the pRolocdata package.

## 6 Semi-supervised machine learning and novelty detection

The PhenoDisco algorithm is a semi-supervised novelty detection method by Breckels et al. (2013) (figure 2). It uses the labelled (i.e. markers, noted  $D_L$ ) and unlabelled (i.e. proteins of unknown localisation, noted  $D_U$ ) sets of the input data. The algorithm is repeated  $N$  times (the `code` argument in the `phenoDisco` function). At each iteration, each organelle class  $D_L^i$  and the unlabelled complement are clustered using Gaussian mixture modelling. While unlabelled members that systematically cluster with  $D_L^i$  and pass outlier detection are labelled as new putative members of class  $i$ , any example of  $D_U$  which are not merged with any any of the  $D_L^i$  and are consistently clustered together throughout the  $N$  iterations are considered members of a new phenotype.





**Figure 2:** The PhenoDisco iterative algorithm.

## Session information

All software and respective versions used to produce this document are listed below.

- R version 3.0.2 Patched (2013-12-18 r64488), x86\_64-unknown-linux-gnu
- Locale: LC\_CTYPE=en\_US.UTF-8, LC\_NUMERIC=C, LC\_TIME=en\_US.UTF-8, LC\_COLLATE=C, LC\_MONETARY=en\_US.UTF-8, LC\_MESSAGES=en\_US.UTF-8, LC\_PAPER=en\_US.UTF-8, LC\_NAME=C, LC\_ADDRESS=C, LC\_TELEPHONE=C, LC\_MEASUREMENT=en\_US.UTF-8, LC\_IDENTIFICATION=C
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, utils
- Other packages: AnnotationDbi 1.24.0, Biobase 2.22.0, BiocGenerics 0.8.0, MASS 7.3-29, MLInterfaces 1.42.0, MSnbase 1.10.3, Rcpp 0.11.0, annotate 1.40.0, class 7.3-9, cluster 1.14.4, codetools 0.2-8, e1071 1.6-2, genefilter 1.44.0, ggplot2 0.9.3.1, knitr 1.5, mzR 1.8.1, pRoloc 1.2.1, pRolocdata 1.0.3, rda 1.0.2-2, rpart 4.1-5, sfsmisc 1.0-25, xtable 1.7-1
- Loaded via a namespace (and not attached): BioInstaller 1.12.0, DBI 0.2-7, FNN 1.1, IRanges 1.20.6, KernSmooth 2.23-10, MSBVAR 0.7-2, Matrix 1.1-2, RColorBrewer 1.0-5, RSQLite 0.11.4, XML 3.98-1.1, affy 1.40.0, affyio 1.30.0, car 2.0-19, caret 6.0-22, coda 0.16-1, colorspace 1.2-4, dichromat 2.0-0, digest 0.6.4, evaluate 0.5.1, foreach 1.4.1, formatR 0.10, gdata 2.13.2, grid 3.0.2, gtable 0.1.2, gtools 3.2.1, highr 0.3, impute 1.36.0, iterators 1.0.6, kernlab 0.9-19, labeling 0.2, lattice 0.20-24, limma 3.18.10, lpSolve 5.6.7, mboost 2.2-3, mclust 4.2, munsell 0.4.2, nnet 7.3-7, plyr 1.8, preprocessCore 1.24.0, proto 0.3-10, proxy 0.4-12, randomForest 4.6-7, reshape2 1.2.2, sampling 2.6, scales 0.2.3, splines 3.0.2, stats4 3.0.2, stringr 0.6.2, survival 2.37-7, tools 3.0.2, vsn 3.30.0, zlibbioc 1.8.0

## References

- Jens S. Andersen, Christopher J. Wilkinson, Thibault Mayor, Peter Mortensen, Erich A. Nigg, and Matthias Mann. Proteomic characterization of the human centrosome by protein correlation profiling. *Nature*, 426(6966):570–574, Dec 2003. doi: 10.1038/nature02166. URL <http://dx.doi.org/10.1038/nature02166>.
- Lisa M Breckels, Laurent Gatto, Andy Christoforou, Arnoud J Groen, Kathryn S Lilley, and Matthew W B Trotter. The effect of organelle discovery upon sub-cellular protein localisation. *J Proteomics*, Mar 2013. doi: 10.1016/j.jprot.2013.02.019.
- Vince Carey, Robert Gentleman, Jess Mar, , contributions from Jason Vertrees, and Laurent Gatto. *MLInterfaces: Uniform interfaces to R machine learning procedures for data in Bioconductor containers*. R package version 1.41.1.
- Nicolas Courty, Thomas Burger, and Johann Laurent. Perturbo: A new classification algorithm based on the spectrum perturbations of the laplace-beltrami operator.

In Dimitrios Gunopulos, Thomas Hofmann, Donato Malerba, and Michalis Vazirgiannis, editors, *in the proceedings of ECML/PKDD (1)*, volume 6911 of *Lecture Notes in Computer Science*, pages 359–374. Springer, 2011.

Tom P. J. Dunkley, Svenja Hester, Ian P. Shadforth, John Runions, Thilo Weimar, Sally L. Hanton, Julian L. Griffin, Conrad Bessant, Federica Brandizzi, Chris Hawes, Rod B. Watson, Paul Dupree, and Kathryn S. Lilley. Mapping the arabidopsis organelle proteome. *Proc Natl Acad Sci USA*, 103(17):6518–6523, Apr 2006. doi: 10.1073/pnas.0506958103. URL <http://dx.doi.org/10.1073/pnas.0506958103>.

N Nikolovski, D Rubtsov, M P Segura, G P Miles, T J Stevens, T P Dunkley, S Munro, K S Lilley, and P Dupree. Putative glycosyltransferases and other plant golgi apparatus proteins are revealed by LOPIT proteomics. *Plant Physiol*, 160(2):1037–51, Oct 2012. doi: 10.1104/pp.112.204263.

S Ohta, J C Bukowski-Wills, L Sanchez-Pulido, Fde L Alves, L Wood, Z A Chen, M Platani, L Fischer, D F Hudson, C P Ponting, T Fukagawa, W C Earnshaw, and J Rappsilber. The protein composition of mitotic chromosomes determined using multiclassifier combinatorial proteomics. *Cell*, 142(5):810–21, Sep 2010. doi: 10.1016/j.cell.2010.07.047.

Denise J Tan, Heidi Dvinge, Andy Christoforou, Paul Bertone, Alfonso A Martinez, and Kathryn S Lilley. Mapping organelle proteins and protein complexes in drosophila melanogaster. *J Proteome Res*, 8(6):2667–78, Jun 2009. doi: 10.1021/pr800866n.

M Tardif, A Atteia, M Specht, G Cogne, N Rolland, S Brugire, M Hippler, M Ferro, C Bruley, G Peltier, O Vallon, and L Cournac. Predalگو: a new subcellular localization prediction tool dedicated to green algae. *Mol Biol Evol*, 29(12):3625–39, Dec 2012. doi: 10.1093/molbev/mss178.

Matthew W. B. Trotter, Pawel G. Sadowski, Tom P. J. Dunkley, Arnoud J. Groen, and Kathryn S. Lilley. Improved sub-cellular resolution via simultaneous analysis of organelle proteomics data across varied experimental conditions. *PROTEOMICS*, 10(23):4213–4219, 2010. ISSN 1615-9861. doi: 10.1002/pmic.201000359. URL <http://dx.doi.org/10.1002/pmic.201000359>.

Sebastian Wiese, Thomas Gronemeyer, Rob Ofman, Markus Kunze, Cludia P. Grou, Jos A. Almeida, Martin Eisenacher, Christian Stephan, Heiko Hayen, Lukas Schollenberger, Thomas Korosec, Hans R. Waterham, Wolfgang Schliebs, Ralf Erdmann, Johannes Berger, Helmut E. Meyer, Wilhelm Just, Jorge E. Azevedo, Ronald J. A. Wanders, and Bettina Warscheid. Proteomics characterization of mouse kidney peroxisomes by tandem mass spectrometry and protein correlation profiling. *Mol Cell Proteomics*, 6(12):2045–2057, Dec 2007. doi: 10.1074/mcp.M700169-MCP200. URL <http://dx.doi.org/10.1074/mcp.M700169-MCP200>.