

# A transfer learning algorithm for spatial proteomics

Lisa M. Breckels and Laurent Gatto\*  
Computational Proteomics Unit  
University of Cambridge, UK

April 16, 2015

---

## Abstract

This vignette illustrates the application of a *transfer learning* algorithm to assign proteins to sub-cellular localisations. The *knntlClassification* algorithm combines *primary* experimental spatial proteomics data (LOPIT, PCP, etc.) and an *auxiliary* data set (for example binary data based on Gene Ontology terms) to improve the sub-cellular assignment given an optimal combination of these data sources.

*Keywords:* Bioinformatics, organelle, spatial proteomics, machine learning, transfer learning

---

---

\*[lg390@cam.ac.uk](mailto:lg390@cam.ac.uk)

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Preparing the auxiliary data</b>	<b>3</b>
2.1	The Gene Ontology . . . . .	3
2.1.1	Preparing the query parameters . . . . .	4
2.1.2	Preparing the auxiliary data from the GO ontology . . . . .	5
2.1.3	A note on reproducibility . . . . .	6
2.2	The Human Protein Atlas . . . . .	6
<b>3</b>	<b>Optimal weights</b>	<b>8</b>
3.1	Choosing weights . . . . .	12
<b>4</b>	<b>Applying best theta weights</b>	<b>14</b>
<b>5</b>	<b>Conclusions</b>	<b>14</b>

## 1 Introduction

---

Our main data source to study protein sub-cellular localisation are high-throughput mass spectrometry-based experiments such as LOPIT, PCP and similar designs (see [1] for an general introduction). Recent optimised experiments result in high quality data enabling the identification of over 6000 proteins and discriminate numerous sub-cellular and sub-organellar niches [2]. Supervised and semi-supervised machine learning algorithms can be applied to assign thousands of proteins to annotated sub-cellular niches [3, 4] (see also the *pRoloc-tutorial* vignette). These data constitute our main source for protein localisation and are termed thereafter *primary* data.

There are other sources of data about sub-cellular localisation of proteins, such as the Gene Ontology [5] (in particular the cellular compartment name space), quantitative features derived from protein sequences (such as pseudo amino acid composition) or the Human Protein Atlas [6] to cite a few. These data, while not optimised to a specific system at hand and, in the case of annotation feature, not as reliable as our experimental data, constitute an invaluable, often plentiful source of *auxiliary* information.

The aim of a *transfer learning* algorithm is to combine different sources of data to improve overall classification. In particular, the goal is to support/complement the primary target domain (experimental data) with auxiliary data (annotation) features without compromising the integrity of our primary data. In this vignette, we describe the application of transfer learning algorithms for the localisation of proteins from the *pRoloc* package [2].

```
> library("pRoloc")
```

## 2 Preparing the auxiliary data

---

### 2.1 The Gene Ontology

The auxiliary data is prepared from the primary data's features. All the GO terms associated to these features are retrieved and used to create a binary matrix where a one (zero) at position  $(i, j)$  indicates that term  $j$  has (not) been used to annotate feature  $i$ .

The GO terms are retrieved from an appropriate repository using the *biomaRt* package. The specific Biomart repository and query will depend on the species under study and the type of features. The first step is to prepare annotation parameters that will enable to perform the query. The *pRoloc* package provides a dedicated infrastructure to set up the query to the annotation resource and prepare the GO data for subsequent analyses. This infrastructure is composed of:

1. define the annotation parameters based on the species and feature types;
2. query the resource defined in (1) to retrieve relevant terms and use the terms to prepare the auxiliary data.

We will demonstrate these steps using a LOPIT experiment on Human Embryonic Kidney (HEK293T) fibroblast cells [3], available and documented in the *pRolocdata* experiment package as andy2011.

```
> library("pRolocdata")
> data(andy2011)
```

### 2.1.1 Preparing the query parameters

The query parameters are stored as *AnnotationParams* objects that are created with the `setAnnotationParams` function. The function will present a first menu with 215. Once the species has been selected, a set of possible identifier types is displayed.




Figure 1: Selecting species (left) and feature type (right) to create an *AnnotationParams* instance for the human *andy2011* data.

It is also possible to pass patterns<sup>1</sup> to match against the species ("Homo sapiens") and feature type ("UniProt/Swissprot ID").

```
> ap <- setAnnotationParams(inputs =
+                               c("Homo sapiens",
+                               "UniProt/Swissprot ID"))
```

*Using species Homo sapiens genes (GRCh38)*

*Using feature type UniProt/Swissprot ID(s) [e.g. Q9H3N8]*

Connecting to BioMart...

```
> ap
```

Object of class "AnnotationParams"

Using the 'ensembl' BioMart database

Using the 'hsapiens\_gene\_ensembl' dataset

Using 'uniprot\_swissprot' as filter

Created on Thu Apr 16 20:51:23 2015

The `setAnnotationParams` function automatically sets the annotation parameters globally so that the `ap` object does not need to be explicitly set later on. The default parameters can be retrieved with

<sup>1</sup>These patterns must match uniquely or an error will be thrown.

```
getAnnotationParams.
```

### 2.1.2 Preparing the auxiliary data from the GO ontology

The feature names of the andy2011 data are UniProt identifiers, as defined in the ap accession parameters.

```
> data(andy2011)
> head(featureNames(andy2011))

[1] "000767" "P51648" "Q2TAA5" "Q9UKV5" "Q12797"
[6] "P16615"
```

The makeGoSet function takes an *MSnSet* class (from which the feature names will be extracted) or, directly a vector of characters containing the feature names of interest to retrieve the associated GO terms and construct an auxiliary *MSnSet*. By default, it downloads *cellular component* terms and does not do any filtering on the terms evidence codes (see the makeGoSet manual for details). Unless passed as argument, the default, globally set *AnnotationParams* are used to define the Biomart server and the query<sup>2</sup>.

```
> andygoset <- makeGoSet(andy2011)
> andygoset

MSnSet (storageMode: lockedEnvironment)
assayData: 1371 features, 673 samples
  element names: exprs
protocolData: none
phenoData: none
featureData
  featureNames: 000767 P51648 ... 075312 (1371
    total)
  fvarLabels: Accession.No. Protein.Description
    ... UniProtKB.entry.name (10 total)
  fvarMetadata: labelDescription
experimentData: use 'experimentData(object)'
Annotation:
- - - Processing information - - -
Constructed GO set using cellular_component namespace: Thu Apr 16 20:51:36 2015
MSnbase version: 1.16.0

> exprs(andygoset)[1:7, 1:4]

      GO:0005783 GO:0005789 GO:0016020 GO:0016021
000767          1          1          1          1
P51648          1          1          0          1
Q2TAA5          1          1          1          1
```

<sup>2</sup>The annotation parameters could also be passed explicitly through the params argument.

Q9UKV5	1	1	1	1
Q12797	1	1	1	1
P16615	1	1	1	1
Q96SQ9	1	1	0	0

We now have a primary data set, composed of 1371 protein quantitative profiles for 8 fractions along the density gradient and an auxiliary data set for 673 cellular compartment GO terms for the same 1371 features.

### 2.1.3 A note on reproducibility

The generation of the auxiliary data relies on specific Biomart server *Mart* instances in the *AnnotationParams* class and the actual query to the server to obtain the GO terms associated with the features. The utilisation of online servers, which undergo regular updates, does not guarantee reproducibility of feature/term association over time. It is recommended to save and store the *AnnotationParams* and auxiliary *MSnSet* instances. Alternatively, it is possible to use other Bioconductor infrastructure, such as specific organism annotations and the *GO.db* package to use specific versioned (and thus traceable) annotations.

## 2.2 The Human Protein Atlas

The feature names of our LOPIT experiment are UniProt identifiers, while the Human Protein Atlas uses Ensembl gene identifiers. This first code chunk matches both identifier types using the UniProt Biomart server.

```
> fvarLabels(andy2011)[1] <- "accession" ## for left_join matching
> ## convert protein accession numbers to ensembl gene identifiers
> library("biomaRt")
> uniprot <- useMart("unimart", dataset = "uniprot")
> filter <- "accession"
> attrib <- c("name", "accession", "ensembl_id")
> bm <- getBM(attributes = attrib,
+             filters = filter,
+             values = fData(andy2011)[, "accession"],
+             mart = uniprot)
> ## HPA data
> library("hpar")
```

*This is hpar 1.10.0. For more information, please type '?hpar' or 'vignette('hpar')'.*

```
> getHpaVersion()
[1] "Protein Atlas version 13"
```

```

> getHpaDate()
[1] "2014.11.06"

> setHparOptions(hpdata = "SubcellularLoc")
> hpa <- getHpa(bm$ensembl_id)
> hpa$Reliability <- droplevels(hpa$Reliability)
> colnames(hpa)[1] <- "ensembl_id"
>
> library("dplyr")
> hpa <- left_join(hpa, bm)

Joining by: "ensembl_id"

Warning in left_join_impl(x, y, by$x, by$y): joining character vector and factor,
coercing into character vector

> hpa <- hpa[!duplicated(hpa$accession), ]
>
> ## match HPA/LOPIT
> fd <- left_join(fData(andy2011), hpa)

Joining by: "accession"

Warning in left_join_impl(x, y, by$x, by$y): joining character vector and factor,
coercing into character vector

> rownames(fd) <- featureNames(andy2011)
> fData(andy2011) <- fd
> stopifnot(validObject(andy2011))
>
> ## Let's get rid of features without any hpa data
> lopit <- andy2011[!is.na(fData(andy2011)$Main.location), ]

```

Below, we deparse the multiple ';' delimited locations contained in the Human Protein sub-cellular Atlas, create the auxiliary binary data matrix (only localisations with reliability equal to *Supportive* are considered; *Uncertain* assignments are ignored) and filter proteins without any localisation data.

```

> ## HPA localisation
> hpalocs <- c(as.character(fData(lopit)$Main.location),
+             as.character(fData(lopit)$Other.location))
> hpalocs <- hpalocs[!is.na(hpalocs)]
> hpalocs <- unique(unlist(strsplit(hpalocs, ";")))
>
> makeHpaSet <- function(x, score2, locs = hpalocs) {
+   hpamat <- matrix(0, ncol = length(locs), nrow = nrow(x))
+   colnames(hpamat) <- locs
+   rownames(hpamat) <- featureNames(x)
+   for (i in 1:nrow(hpamat)) {

```

```

+     loc <- unlist(strsplit(as.character(fData(x)[i, "Main.location"]), ";"))
+     loc2 <- unlist(strsplit(as.character(fData(x)[i, "Other.location"]), ";"))
+     score <- score2[fData(x)[i, "Reliability"]]
+     hpamat[i, loc] <- score
+     hpamat[i, loc2] <- score
+   }
+   new("MSnSet", exprs = hpamat,
+       featureData = featureData(lopit))
+ }
>
> hpaset <- makeHpaSet(lopit,
+                      score2 = c(Supportive = 1, Uncertain = 0))
> hpaset <- filterZeroRows(hpaset)

```

*Removing 317 columns with only 0s.*

```

> dim(hpaset)
[1] 670 18
> exprs(hpaset)[c(1, 6, 200), 1:3]
      Endoplasmic reticulum Cytoplasm Vesicles
000767                1         0         0
095302                0         0         1
P06493                0         1         0

```

### 3 Optimal weights

The weighted nearest neighbours transfer learning algorithm estimates optimal weights for the different data sources and the spatial niches described for the data at hand with the `knnt10optimisation` function. For instance, for the human data modelled by the `andy2011` and `andygoaset` objects<sup>3</sup> and the 10 annotated sub-cellular localisations (Golgi, Mitochondrion, PM, Lysosome, Cytosol, Cytosol/Nucleus, Nucleus, Ribosome 60S, Ribosome 40S and ER), we want to know how to optimally combine primary and auxiliary data. If we look at figure 2, that illustrates the experimental separation of the 10 spatial classes on a principal component plot, we see that some organelles such as the mitochondrion or the cytosol and cytosol/nucleus are well resolved, while others such as the Golgi or the ER are less so. In this experiment, the former classes are not expected to benefit from another data source, while the latter should benefit from additional information.

Let's define a set of three possible weights: 0, 0.5 and 1. A weight of 1 indicates that the final results rely exclusively on the experimental data and ignore completely the auxiliary data. A weight of 0 represents the opposite situation, where the primary data is ignored and only the auxiliary data is considered. A weight of 0.5 indicates that each data source will contribute equally to the final results. It is the

<sup>3</sup>We will use the sub-cellular markers defined in the `markers.t1` feature variable, instead of the default `markers`.






Figure 2: PCA plot of andy2011. The multivariate protein profiles are summarised along the two first principal components. Proteins of unknown localisation are represented by empty grey points. Protein markers, which are well-known residents of specific sub-cellular niches are colour-coded and form clusters on the figure.

algorithm's optimisation step task to identify the optimal combination of class-specific weights for a given primary and auxiliary data pair. The optimisation process can be quite time consuming for many weights and many sub-cellular classes, as all combinations (there are  $\text{number of classes}^{\text{number of weights}}$  possibilities; see below). One would generally defined more weights (for example 0, 0.25, 0.5, 0.75, 1 or 0, 0.33, 0.67, 1) to explore more fine-grained integration opportunities. The possible weight combinations can be calculated with the `thetas` function:

- 3 classes, 3 weights

```
> head(thetas(3, by = 0.5))
      [,1] [,2] [,3]
[1,]    0  0.0  0.0
[2,]    0  0.0  0.5
[3,]    0  0.0  1.0
[4,]    0  0.5  0.0
[5,]    0  0.5  0.5
[6,]    0  0.5  1.0
> dim(thetas(3, by = 0.5))
Weights:
(0, 0.5, 1)
[1] 27 3
```

- 5 classes, 4 weights

```
> dim(thetas(5, length.out = 4))
Weights:
(0, 0.3333333333333333, 0.6666666666666667, 1)
[1] 1024 5
```

- for the human andy2011 data, considering 4 weights, there are very many combinations:

```
> ## marker classes for andy2011
> m <- unique(fData(andy2011)$markers.t1)
> m <- m[m != "unknown"]
> th <- thetas(length(m), length.out=4)
Weights:
(0, 0.3333333333333333, 0.6666666666666667, 1)
> dim(th)
[1] 1048576 10
```

The actual combination of weights to be tested can be defined in multiple ways: by passing a weights matrix explicitly (as those generated with `thetas` above) through the `th` argument; or by defining the increment between weights using `by`; or by specifying the number of weights to be used through the `length.out` argument.

Considering the sub-cellular resolution for this experiment, we would anticipate that the mitochondrion, the cytosol and the cytosol/nucleus classes would get high weights, while the ER and Golgi would be assigned lower weights.

As we use a nearest neighbour classifier, we also need to know how many neighbours to consider when classifying a protein of unknown localisation. The `knnOptimisation` function (see the *pRoloc-tutorial* vignette and the functions manual page) can be run on the primary and auxiliary data sources independently to estimate the best  $k_P$  and  $k_A$  values. Here, based on `knnOptimisation`, we use 3 and 3, for  $k_P$  and  $k_A$  respectively.

Finally, to assess the validity of the weight selection, it should be repeated a certain number of times (default value is 50). As the weight optimisation can become very time consuming for a wide range of weights and many target classes, we would recommend to start with a lower number of iterations,

pre-analyse the results, proceed with further iterations and eventually combine the optimisation results data with the `combineThetaRegRes` function before proceeding with the selection of best weights.

```
> topt <- knnt1Optimisation(andy2011, andygoset,
+                           th = th,
+                           k = c(3, 3),
+                           fcol = "markers.tl",
+                           times = 50)
```

The above code chunk would take too much time to be executed in the frame of this vignette. Below, we pass a very small subset of theta matrix to minimise the computation time. The `knnt1Optimisation` function supports parallelised execution using various backends thanks to the [BiocParallel](#) package; an appropriate backend will be defined automatically according to the underlying architecture and user-defined backends can be defined through the `BPPARAM` argument<sup>4</sup>. Also, in the interest of time, the weights optimisation is repeated only 5 times below.

```
> set.seed(1)
> i <- sample(nrow(th), 12)
> topt <- knnt1Optimisation(andy2011, andygoset,
+                           th = th[i, ],
+                           k = c(3, 3),
+                           fcol = "markers.tl",
+                           times = 5)
```

*Removing 356 columns with only 0s.*

*Note: vector will be ordered according to classes: Cytosol Cytosol/Nucleus ER Golgi Lysosome Mitochondrion Nucleus PM Ribosome 40S Ribosome 60S (as names are not explicitly defined)*

```
> topt
```

```
Object of class "ThetaRegRes"
```

```
Algorithm: theta
```

```
Theta hyper-parameters:
```

```
weights: 0 0.3333333 0.6666667 1
```

```
k: 3 3
```

```
nrow: 12
```

```
Design:
```

```
Replication: 5 x 5-fold X-validation
```

```
Partitioning: 0.2/0.8 (test/train)
```

```
Results
```

```
macro F1:
```

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.8974	0.9085	0.9174	0.9193	0.9251	0.9483

```
best theta:
```

<sup>4</sup>Large scale applications of this algorithms (*fixme: add ref*) were run on a cluster using an MPI backend defined with `SnowParams(256, type="MPI")`.

```

      Cytosol Cytosol.Nucleus ER Golgi Lysosome
weight:0      0              0 5      0      4
weight:0.33    0              0 0      5      0
weight:0.67    4              4 0      0      0
weight:1      1              1 0      0      1

      Mitochondrion Nucleus PM Ribosome.40S
weight:0          4          0 0          0
weight:0.33        1          1 5          4
weight:0.67        0          0 0          0
weight:1          0          4 0          1

      Ribosome.60S
weight:0          4
weight:0.33        0
weight:0.67        1
weight:1          0
Use getWarnings() to see warnings.

```

The optimisation is performed on the labelled marker examples only. When removing unlabelled non-marker proteins (the `unknowns`), some auxiliary GO columns end up containing only 0 (the GO-protein association was only observed in non-marker proteins), which are temporarily removed.

The `topt` result stores all the result from the optimisation step, and in particular the observed theta weights, which can be directly plotted as shown on figure 3. These *bubble* plots give the proportion of best weights for each marker class that was observed during the optimisation phase. We see that the mitochondrion, the cytosol and cytosol/nucleus classes predominantly are scored with height weights (2/3 and 1), consistent with high reliability of the primary data. The Golgi and the ribosomal clusters (and to a lesser extend the ER) favour smaller scores, indicating a substantial benefit of the auxiliary data.

### 3.1 Choosing weights

A set of best weights must be chosen and applied to the classification of the unlabelled proteins (formally annotated as `unknown`). These can be defined manually, based on the pattern observed in the weights *bubble* plot (figure 3), or automatically extracted with the `getParams` method<sup>5</sup>. See `?getParams` for details and the `favourPrimary` function, if it is desirable to systematically favour the primary data (i.e. high weights) when different weight combinations perform equally well.

```

> getParams(topt)

      Cytosol Cytosol/Nucleus      ER
0.6666667    0.6666667    0.0000000
      Golgi      Lysosome Mitochondrion
0.3333333    0.0000000    0.0000000

```

<sup>5</sup>Note that the scores extracted here are based on the random subset of weights.




Figure 3: Results obtained from an extensive optimisation on the primary andy2011 and auxiliary andygoset data sets, as produced by `plot(topt)`. This figure is not the result for the previous code chunk, where only a random subset of 10 candidate weights have been tested.

Nucleus	PM	Ribosome 40S
1.0000000	0.3333333	0.3333333
Ribosome 60S		
0.0000000		

We provide the best parameters for the extensive parameter optimisation search, as provided by `getParams`:

```
> (bw <- experimentData(andy2011)@other$knnt1$thetas)
      Cytosol Cytosol/Nucleus      ER
0.6666667    0.6666667    0.3333333
      Golgi      Lysosome Mitochondrion
0.3333333    0.6666667    0.6666667
      Nucleus      PM Ribosome 40S
```

0.3333333	0.3333333	0.0000000
Ribosome 60S		
0.3333333		

## 4 Applying best theta weights

To apply our best weights and learn from the auxiliary data accordingly when classifying the unlabelled proteins to one of the sub-cellular niches considered in `markers.tl` (as displayed on figure 2), we pass the primary and auxiliary data sets, best weights, best  $k$ 's (and, on our case the marker's feature variable we want to use, default would be `markers`) to the `knnt1Classification` function.

```
> andy2011 <- knnt1Classification(andy2011, andygoset,
+                               bestTheta = bw,
+                               k = c(3, 3),
+                               fcol = "markers.tl")
```

This will generate a new instance of class *MSnSet*, identical to the primary data, including the classification results and classifications scores of the transfer learning classification algorithm (as `knnt1` and `knnt1.scores` feature variables respectively). Below, we extract the former with the `getPrediction` function and plot the results of the classification.

```
> getPredictions(andy2011, fcol = "knnt1")
```

ans

Cytosol	Cytosol/Nucleus	ER
284	48	202
Golgi	Lysosome	Mitochondrion
66	88	270
Nucleus	PM	Ribosome 40S
143	194	18
Ribosome 60S		
58		

Please read the *pRoloc-tutorial* vignette, and in particular the classification section, for more details on how to proceed with exploration the classification results and classification scores.

## 5 Conclusions

This vignette describes the application of a weighted  $k$ -nearest neighbour transfer learning algorithm and its application to the sub-cellular localisation prediction of proteins using quantitative proteomics data as primary data and Gene Ontology-derived binary data as auxiliary data source. The algorithm can be used with various data sources (we show how to compile binary data from the Human Protein

Atlas in section 2.2) and have successfully applied the algorithm [2] on third-party quantitative auxiliary data.

## Session information

---

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

- R version 3.2.0 RC (2015-04-08 r68161), x86\_64-apple-darwin13.4.0
- Locale: C/en\_US.UTF-8/en\_US.UTF-8/C/en\_US.UTF-8/en\_US.UTF-8
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, stats4, utils
- Other packages: AnnotationDbi 1.30.0, Biobase 2.28.0, BiocGenerics 0.14.0, BiocParallel 1.2.0, GenomInfoDb 1.4.0, IRanges 2.2.0, MLInterfaces 1.48.0, MSnbase 1.16.0, ProtGenerics 1.0.0, Rcpp 0.11.5, S4Vectors 0.6.0, XML 3.98-1.1, annotate 1.46.0, biomaRt 2.24.0, class 7.3-12, cluster 2.0.1, dplyr 0.4.1, hpar 1.10.0, knitr 1.9, mzR 2.2.0, pRoloc 1.8.0, pRolocdata 1.5.8, xtable 1.7-4
- Loaded via a namespace (and not attached): BiocInstaller 1.18.1, BiocStyle 1.6.0, BradleyTerry2 1.0-6, DBI 0.3.1, FNN 1.1, MALDIquant 1.11, MASS 7.3-40, Matrix 1.2-0, RColorBrewer 1.1-2, RCurl 1.95-4.5, RSQLite 1.0.0, SparseM 1.6, affy 1.46.0, affyio 1.36.0, assertthat 0.1, bitops 1.0-6, brglm 0.5-9, car 2.0-25, caret 6.0-41, codetools 0.2-11, colorspace 1.2-6, digest 0.6.8, doParallel 1.0.8, e1071 1.6-4, evaluate 0.6, foreach 1.4.2, formatR 1.1, futile.logger 1.4, futile.options 1.0.0, gdata 2.13.3, genefilter 1.50.0, ggplot2 1.0.1, grid 3.2.0, gtable 0.1.2, gtools 3.4.2, highr 0.4.1, impute 1.42.0, iterators 1.0.7, kernlab 0.9-20, lambda.r 1.1.7, lattice 0.20-31, limma 3.24.0, lme4 1.1-7, lpSolve 5.6.11, magrittr 1.5, mclust 5.0.0, mgcv 1.8-6, minqa 1.2.4, munsell 0.4.2, mvtnorm 1.0-2, mzID 1.6.0, nlme 3.1-120, nloptr 1.0.4, nnet 7.3-9, pbkrtest 0.4-2, pcaMethods 1.58.0, pls 2.4-3, plyr 1.8.1, preprocessCore 1.30.0, proto 0.3-10, proxy 0.4-14, quantreg 5.11, randomForest 4.6-10, rda 1.0.2-2, reshape2 1.4.1, rpart 4.1-9, sampling 2.6, scales 0.2.4, sfsmisc 1.0-27, snow 0.3-13, splines 3.2.0, stringr 0.6.2, survival 2.38-1, tools 3.2.0, vsn 3.36.0, zlibbioc 1.14.0

## References

---

- [1] Laurent Gatto, Juan Antonio Vizcaíno, Henning Hermjakob, Wolfgang Huber, and Kathryn S Lilley. Organelle proteomics experimental designs and analysis. *Proteomics*, 2010. doi:10.1002/pmic.201000244.
- [2] L M Breckels, S Holden, D Wonjar, A Christoforou, A Groen, O Kohlbacher, K S Lilley, and L Gatto. A transfer learning framework for spatial proteomics. *in prep*.
- [3] 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.

- [4] L Gatto, L M Breckels, T Burger, D J Nightingale, A J Groen, C Campbell, N Nikolovski, C M Mulvey, A Christoforou, M Ferro, and K S Lilley. A foundation for reliable spatial proteomics data analysis. *Mol Cell Proteomics*, 13(8):1937–52, Aug 2014. doi:[10.1074/mcp.M113.036350](https://doi.org/10.1074/mcp.M113.036350).
- [5] M Ashburner, C A Ball, J A Blake, D Botstein, H Butler, J M Cherry, A P Davis, K Dolinski, S S Dwight, J T Eppig, M A Harris, D P Hill, L Issel-Tarver, A Kasarskis, S Lewis, J C Matese, J E Richardson, M Ringwald, G M Rubin, and G Sherlock. Gene ontology: tool for the unification of biology. the gene ontology consortium. *Nat Genet*, 25(1):25–9, May 2000. doi:[10.1038/75556](https://doi.org/10.1038/75556).
- [6] Mathias Uhlen, Per Oksvold, Linn Fagerberg, Emma Lundberg, Kalle Jonasson, Mattias Forsberg, Martin Zwahlen, Caroline Kampf, Kenneth Wester, Sophia Hober, Henrik Wernerus, Lisa Björling, and Fredrik Ponten. Towards a knowledge-based Human Protein Atlas. *Nature biotechnology*, 28(12):1248–1250, December 2010. URL: <http://dx.doi.org/10.1038/nbt1210-1248>, doi:[10.1038/nbt1210-1248](https://doi.org/10.1038/nbt1210-1248).



```

> setStockcol(paste0(getStockcol(), "80"))
> ptsze <- exp(fData(andy2011)$knnt1.scores) - 1
> plot2D(andy2011, fcol = "knnt1", cex = ptsze)
> setStockcol(NULL)
> addLegend(andy2011, where = "topright",
+           fcol = "markers.tl",
+           bty = "n", cex = .7)

```




Figure 4: PCA plot of andy2011 after transfer learning classification. The size of the points is proportional to the classification scores.