

SPAMS: a SPArse Modeling Software, v2.6

Julien Mairal
julien.mairal@m4x.org

Python and R interfaces by Jean-Paul Chieze
Archetypal Analysis by Yuansi Chen
Porting to R-3.x/Python-3.x by Ghislain Durif
and the output of numerous collaborations with
Rodolphe Jenatton, Francis Bach, Jean Ponce,
Guillaume Obozinski, Bin Yu, Guillermo Sapiro,
and Zaid Harchaoui.

February 28, 2017

Contents

1	Introduction	2
2	Installation	4
3	Dictionary Learning and Matrix Factorization Toolbox	4
3.1	Function spams.trainDL	5
3.2	Function spams.trainDL_Memory	9
3.3	Function spams.structTrainDL	12
3.4	Function nmf	17
3.5	Function nnsc	19
3.6	Function spams.archetypalAnalysis	20
4	Sparse Decomposition Toolbox	21
4.1	Function spams.omp	21
4.2	Function spams.ompMask	23
4.3	Function mexRidgeRegression	25
4.4	Function spams.lasso	25
4.5	Function spams.lassoWeighted	28
4.6	Function spams.lassoMask	30
4.7	Function spams.cd	31
4.8	Function spams.somp	33
4.9	Function spams.l1L2BCD	34
4.10	Function spams.sparseProject	36
4.11	Function spams.decompSimplex	39
5	Proximal Toolbox	39
5.1	Regularization Functions	39
5.2	Function spams.proximalFlat	41
5.3	Function spams.proximalTree	45
5.4	Function spams.proximalGraph	50
5.5	Function spams.proximalPathCoding	53
5.6	Function spams.evalPathCoding	53
5.7	Problems Addressed	54

5.7.1	Regression Problems with the Square Loss	54
5.7.2	Classification Problems with the Logistic Loss	54
5.7.3	Multi-class Classification Problems with the Softmax Loss	54
5.7.4	Multi-task Regression Problems with the Square Loss	54
5.7.5	Multi-task Classification Problems with the Logistic Loss	55
5.7.6	Multi-task and Multi-class Classification Problems with the Softmax Loss	55
5.8	Function spams.fistaFlat	55
5.9	Function spams.fistaTree	63
5.10	Function spams.fistaGraph	68
5.11	Function spams.fistaPathCoding	75
5.12	Function solverPoisson	75
5.13	Function mexIncrementalProx	75
5.14	Function mexStochasticProx	75
6	A few Functions for manipulating images	76
6.1	Function mexExtractPatches	76
6.2	Function mexCombinePatches	76
6.3	Function mexConvFistaFlat	76
7	Miscellaneous Functions	76
7.1	Function spams.conjGrad	76
7.2	Function spams.bayer	77
7.3	Function spams.calcAAt	78
7.4	Function spams.calcXAt	79
7.5	Function spams.calcXY	79
7.6	Function spams.calcXYt	80
7.7	Function spams.calcXtY	80
7.8	Function spams.invSym	81
7.9	Function spams.normalize	81
7.10	Function spams.sort	82
7.11	Function displayPatches	83
7.12	Function spams.countPathsDAG	83
7.13	Function spams.removeCyclesGraph	83
7.14	Function spams.countConnexComponents	83
7.15	Function spams.graphOfGroupStruct	83
7.16	Function spams.groupStructOfString	84
7.17	Function spams.readGroupStruct	84
7.18	Function spams.treeOfGroupStruct	85
7.19	Function spams.simpleGroupTree	86
A	Duality Gaps with Fenchel Duality	86
A.0.1	Duality Gaps without Intercepts	87
A.0.2	Duality Gaps with Intercepts	88

1 Introduction

SPAMS (SParse Modeling Software) is an open-source optimization toolbox for sparse estimation with licence GPLv3. It implements algorithms for solving machine learning and signal processing problems involving sparse regularizations.

The library is coded in C++, is compatible with Linux, Mac, and Windows 32bits and 64bits Operating Systems. It is interfaced with Matlab, R and Python (2.x and 3.x), but can be called from any C++ application (by hacking the code a bit).

It requires an implementation of BLAS and LAPACK for performing linear algebra operations. The ones shipped with Matlab and R can be used, but also external libraries such as atlas, the netlib implementation,

or the Intel Math Kernel Library can be used. It also exploits multi-core CPUs when this feature is supported by the compiler, through OpenMP.

The current licence is GPLv3, which is available at <http://www.gnu.org/licenses/gpl.html>. For other licensing possibilities allowing its use in proprietary softwares, please contact the author.

The porting to python3.x is based on <https://aur.archlinux.org/packages/python-spams-svn/>.

Version 2.6 of SPAMS is divided into several “toolboxes” and has a few additional miscellaneous functions:

- The **Dictionary learning and matrix factorization toolbox** contains the online learning technique of [20, 21] and its variants for solving various matrix factorization problems:
 - dictionary Learning for sparse coding;
 - sparse principal component analysis (seen as a sparse matrix factorization problem);
 - non-negative matrix factorization;
 - non-negative sparse coding;
 - dictionary learning with structured sparsity;
 - archetypal analysis [7, 37].
- The **Sparse decomposition toolbox** contains efficient implementations of
 - Orthogonal Matching Pursuit, (or Forward Selection) [35, 27];
 - the LARS/homotopy algorithm [30, 9] (variants for solving Lasso and Elastic-Net problems);
 - a weighted version of LARS;
 - OMP and LARS when data comes with a binary mask;
 - a coordinate-descent algorithm for ℓ_1 -decomposition problems [12, 10, 36];
 - a greedy solver for simultaneous signal approximation as defined in [34, 33] (SOMP);
 - a solver for simultaneous signal approximation with ℓ_1/ℓ_2 -regularization based on block-coordinate descent;
 - a homotopy method for the Fused-Lasso Signal Approximation as defined in [10] with the homotopy method presented in the appendix of [21];
 - a tool for projecting efficiently onto a few convex sets inducing sparsity such as the ℓ_1 -ball using the method of [3, 18, 8], and Elastic-Net or Fused Lasso constraint sets as proposed in the appendix of [21].
 - an active-set algorithm for simplex decomposition problems [37].
- The **Proximal toolbox**: An implementation of proximal methods (ISTA and FISTA [1]) for solving a large class of sparse approximation problems with different combinations of loss and regularizations. One of the main features of this toolbox is to provide a robust stopping criterion based on *duality gaps* to control the quality of the optimization, whenever possible. It also handles sparse feature matrices for large-scale problems. The following regularizations are implemented:
 - Tikhonov regularization (squared ℓ_2 -norm);
 - ℓ_1 -norm, ℓ_2 , ℓ_∞ -norms;
 - Elastic-Net [39];
 - Fused Lasso [32];
 - tree-structured sum of ℓ_2 -norms (see [15, 16]);
 - tree-structured sum of ℓ_∞ -norms (see [15, 16]);
 - general sum of ℓ_∞ -norms (see [22, 23]);
 - mixed ℓ_1/ℓ_2 -norms on matrices [38, 29];
 - mixed ℓ_1/ℓ_∞ -norms on matrices [38, 29];
 - mixed ℓ_1/ℓ_2 -norms on matrices plus ℓ_1 [31, 11];

- mixed ℓ_1/ℓ_∞ -norms on matrices plus ℓ_1 ;
- group-lasso with ℓ_2 or ℓ_∞ -norms;
- group-lasso+ ℓ_1 ;
- multi-task tree-structured sum of ℓ_∞ -norms (see [22, 23]);
- trace norm;
- ℓ_0 pseudo-norm (only with ISTA);
- tree-structured ℓ_0 (only with ISTA);
- rank regularization for matrices (only with ISTA);
- the path-coding penalties of [24].

All of these regularization functions can be used with the following losses

- square loss;
- square loss with missing observations;
- logistic loss, weighted logistic loss;
- multi-class logistic.

This toolbox can also enforce non-negativity constraints, handle intercepts and sparse matrices. There are also a few additional undocumented functionalities, which are available in the source code. For some combinations of loss and regularizers, stochastic and incremental proximal gradient solvers are also implemented [26, 25].

- A few tools for performing linear algebra operations such as a conjugate gradient algorithm, manipulating sparse matrices and graphs.

The toolbox was written by Julien Mairal at INRIA, with the collaboration of Francis Bach (INRIA), Jean Ponce (Ecole Normale Supérieure), Guillermo Sapiro (University of Minnesota), Guillaume Obozinski (INRIA) and Rodolphe Jenatton (INRIA).

R and Python interfaces have been written by Jean-Paul Chieze (INRIA). Porting to R-3.x and Python-3.x was done by Ghislain Durif. The archetypal analysis implementation was written by Yuansi Chen, during an internship at INRIA, with the collaboration of Zaid Harchaoui.

2 Installation

The SPAMS toolbox for R is distributed in source mode only. It should compile on linux and Mac. The installation procedure is described in the file `INSTALL-package`

- Download the tar gzipped file, unpack it.
- Set your environment variable `R_LIBS` appropriately
- Enter directory `spams-R`
- execute


```
R CMD INSTALL --html -l $R_LIBS spams
```

You have the choice of the BLAS library, but the Intel MKL is recommended for the best performance. If you want to add or change libraries, you must modify the file `spams/src/Makevars`. The documentation is available in pdf and html format in the `doc` subdirectory.

3 Dictionary Learning and Matrix Factorization Toolbox

This is the section for dictionary learning and matrix factorization, corresponding to [20, 21].

3.1 Function `spams.trainDL`

This is the main function of the toolbox, implementing the learning algorithms of [21]. Given a training set \mathbf{x}^1, \dots . It aims at solving

$$\min_{\mathbf{D} \in \mathcal{C}} \lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n \min_{\alpha^i} \left(\frac{1}{2} \|\mathbf{x}^i - \mathbf{D}\alpha^i\|_2^2 + \psi(\alpha^i) \right). \quad (1)$$

ψ is a sparsity-inducing regularizer and \mathcal{C} is a constraint set for the dictionary. As shown in [21] and in the help file below, various combinations can be used for ψ and \mathcal{C} for solving different matrix factorization problems. What is more, positivity constraints can be added to α as well. The function admits several modes for choosing the optimization parameters, using the parameter-free strategy proposed in [20], or using the parameters t_0 and ρ presented in [21]. **Note that for problems of a reasonable size, and when ψ is the ℓ_1 -norm, the function `spams.trainDL_Memory` can be faster but uses more memory.**

```
#
# Name: spams.trainDL
#
# Usage: spams.trainDL(X,return_model= FALSE,model= NULL,D = NULL,numThreads = -1,batchsize =
# -1,
#           K= -1,lambda1= NULL,lambda2= 10e-10,iter=-1,t0=1e-5,mode='PENALTY',
#           posAlpha=FALSE,posD=FALSE,expand=FALSE,modeD='L2',whiten=FALSE,clean=TRUE,
#           verbose=TRUE,gamma1=0.,gamma2=0.,rho=1.0,iter_updated=NULL,
#           stochastic_deprecated=FALSE,modeParam=0,batch=FALSE,log_deprecated=FALSE,
#           logName='')
#
# Description:
#   spams.trainDL is an efficient implementation of the
#   dictionary learning technique presented in
#
#   "Online Learning for Matrix Factorization and Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   arXiv:0908.0050
#
#   "Online Dictionary Learning for Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   ICML 2009.
#
#   Note that if you use mode=1 or 2, if the training set has a
#   reasonable size and you have enough memory on your computer, you
#   should use spams.trainDL_Memory instead.
#
#   It addresses the dictionary learning problems
#   1) if mode=0
#   min_{D in C} (1/n) sum_{i=1}^n (1/2) ||x_i - D alpha_i||_2^2 s.t. ...
#                                     ||alpha_i||_1 <= lambda1
#   2) if mode=1
#   min_{D in C} (1/n) sum_{i=1}^n ||alpha_i||_1 s.t. ...
#                                     ||x_i - D alpha_i||_2^2 <= lambda1
#   3) if mode=2
#   min_{D in C} (1/n) sum_{i=1}^n (1/2) ||x_i - D alpha_i||_2^2 + ...
#                                     lambda1 ||alpha_i||_1 + lambda1_2 ||alpha_i||_2^2
#   4) if mode=3, the sparse coding is done with OMP
#   min_{D in C} (1/n) sum_{i=1}^n (1/2) ||x_i - D alpha_i||_2^2 s.t. ...
#                                     ||alpha_i||_0 <= lambda1
#   5) if mode=4, the sparse coding is done with OMP
```

```

#   min_{D in C} (1/n) sum_{i=1}^n ||alpha_i||_0 s.t. ...
#
#           ||x_i - D alpha_i||_2^2 <= lambda1
#   6) if mode=5, the sparse coding is done with OMP
#   min_{D in C} (1/n) sum_{i=1}^n 0.5 ||x_i - D alpha_i||_2^2 + lambda1 ||alpha_i||_0
#
#
#   C is a convex set verifying
#   1) if modeD=0
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 <= 1 }
#   2) if modeD=1
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#           gamma1 ||d_j||_1 <= 1 }
#   3) if modeD=2
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#           gamma1 ||d_j||_1 + gamma2 FL(d_j) <= 1 }
#   4) if modeD=3
#       C={ D in Real^{m x p} s.t. forall j, (1-gamma1) ||d_j||_2^2 + ...
#           gamma1 ||d_j||_1 <= 1 }
#
#   Potentially, n can be very large with this algorithm.
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   return_model:
#       if true the function will return the model
#       as a named list ('A' = A, 'B' = B, 'iter' = n)
#   model: NULL or model (as A,B,iter) to use as initialisation
#   D: (optional) double m x p matrix (dictionary)
#       p is the number of elements in the dictionary
#       When D is not provided, the dictionary is initialized
#       with random elements from the training set.
#   K: (size of the dictionary, optional is D is provided)
#   lambda1: (parameter)
#   lambda2: (optional, by default 0)
#   iter: (number of iterations). If a negative number is
#         provided it will perform the computation during the
#         corresponding number of seconds. For instance iter=-5
#         learns the dictionary during 5 seconds.
#   mode: (optional, see above, by default 2)
#   posAlpha: (optional, adds positivity constraints on the
#             coefficients, false by default, not compatible with
#             mode =3,4)
#   modeD: (optional, see above, by default 0)
#   posD: (optional, adds positivity constraints on the
#         dictionary, false by default, not compatible with
#         modeD=2)
#   gamma1: (optional parameter for modeD >= 1)
#   gamma2: (optional parameter for modeD = 2)
#   batchsize: (optional, size of the minibatch, by default
#             512)
#   iter_updated: (optional, number of BCD iterations for the dictionary
#                update step, by default 1)

```

```

# modeParam: (optimization mode).
#   1) if modeParam=0, the optimization uses the
#       parameter free strategy of the ICML paper
#   2) if modeParam=1, the optimization uses the
#       parameters rho as in arXiv:0908.0050
#   3) if modeParam=2, the optimization uses exponential
#       decay weights with updates of the form
#        $A_{t} \leftarrow \rho A_{t-1} + \alpha_t \alpha_t^T$ 
# rho: (optional) tuning parameter (see paper arXiv:0908.0050)
# t0: (optional) tuning parameter (see paper arXiv:0908.0050)
# clean: (optional, true by default. prunes
#         automatically the dictionary from unused elements).
# verbose: (optional, true by default, increase verbosity)
# numThreads: (optional, number of threads for exploiting
#              multi-core / multi-cpus. By default, it takes the value -1,
#              which automatically selects all the available CPUs/cores).
# expand:    undocumented; modify at your own risks!
# whiten:    undocumented; modify at your own risks!
# stochastic_deprecated:    undocumented; modify at your own risks!
# batch:     undocumented; modify at your own risks!
# log_deprecated:    undocumented; modify at your own risks!
# logName:    undocumented; modify at your own risks!
#
# Output:
# D:    double m x p matrix    (dictionary)
# model: the model as A B iter
#   D <- spams.trainDL(X,return_model = FALSE,...)
#   v <- spams.trainDL(X,return_model = TRUE,...)
#   D <- v[[1]]
#   model <- v[[2]]
#   A = model[['A']]
#   B = model[['B']]
#   iter = model[['iter']]
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
I = readPNG(paste(.imagesDir, 'boat.png', sep= '/'))
if (length(dim(I)) == 3) {
  A = matrix(I,nrow = nrow(I),ncol = 3 * ncol(I))
} else {
  A = I
}

```

```

m = 8;n = 8;
X = spams.im2col_sliding(A,m,n)

X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)

lambda1 = 0.15

##### FIRST EXPERIMENT #####
tic = proc.time()
D <- spams.trainDL(X,K = 100,lambda1 = lambda1, numThreads = 4, batchsize = 400,iter = 1000)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)

#### SECOND EXPERIMENT ####
.printf("***** SECOND EXPERIMENT *****\n")

X1 = X[,1:as.integer(ncol(X) / 2 )]
X2 = X[,as.integer(ncol(X) / 2):ncol(X)]

tic = proc.time()
res = spams.trainDL(X1,return_model = TRUE,K = 100,lambda1 = lambda1, numThreads = 4, batchsize
    = 400,iter = 500)
tac = proc.time()
D = res[[1]]
model = res[[2]]
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

                                # Then reuse the learned model to retrain a few
                                # iterations more.

tic = proc.time()
res = spams.trainDL(X2,model = model,return_model = TRUE,D = D,K = 100,lambda1 = lambda1,
    numThreads = 4, batchsize = 400,iter = 500)
tac = proc.time()
D = res[[1]]
model = res[[2]]
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

##### THIRD & FOURTH EXPERIMENT #####
                                # let us add sparsity to the dictionary itself

.printf('***** THIRD EXPERIMENT *****\n')

tic = proc.time()

```



```

D = spams.trainDL(X,K = 100,lambda1 = lambda1, numThreads = 4, batchSize = 400,iter = 1000,
    modeParam = 0,gamma1 = 0.3,modeD = 'L1L2')
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

.printf('***** FOURTH EXPERIMENT *****\n')

tic = proc.time()
D = spams.trainDL(X,K = 100,lambda1 = lambda1, numThreads = 4, batchSize = 400,iter = 1000,
    modeParam = 0,gamma1 = 0.3,modeD = 'L1L2MU')
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

```

3.2 Function spams.trainDL_Memory

Memory-consuming version of spams.trainDL. This function is well adapted to small/medium-size problems: It requires storing all the coefficients α and is therefore impractical for very large datasets. However, in many situations, one can afford this memory cost and it is better to use this method, which is faster than spams.trainDL. Note that unlike spams.trainDL this function does not allow warm-restart.

```

#
# Name: spams.trainDL_Memory
#
# Usage: spams.trainDL_Memory(X,D = NULL,numThreads = -1,batchsize = -1,K= -1,lambda1= NULL,
    iter=-1,
#         t0=1e-5,mode='PENALTY',posD=FALSE,expand=FALSE,modeD='L2',
#         whiten=FALSE,clean=TRUE,gamma1=0.,gamma2=0.,rho=1.0,iter_updated=1,
#         stochastic_deprecated=FALSE,modeParam=0,batch=FALSE,
#         log_deprecated=FALSE,logName='')
#
# Description:
#   spams.trainDL_Memory is an efficient but memory consuming
#   variant of the dictionary learning technique presented in
#
#   "Online Learning for Matrix Factorization and Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   arXiv:0908.0050
#
#   "Online Dictionary Learning for Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   ICML 2009.
#
#   Contrary to the approaches above, the algorithm here
#   does require to store all the coefficients from all the training
#   signals. For this reason this variant can not be used with large
#   training sets, but is more efficient than the regular online
#   approach for training sets of reasonable size.
#
#   It addresses the dictionary learning problems

```

```

#     1) if mode=1
#     min_{D in C} (1/n) sum_{i=1}^n ||alpha_i||_1 s.t. ...
#                                     ||x_i-Dalpha_i||_2^2 <= lambda1
#
#     2) if mode=2
#     min_{D in C} (1/n) sum_{i=1}^n (1/2)||x_i-Dalpha_i||_2^2 + ...
#                                     lambda1||alpha_i||_1
#
# C is a convex set verifying
#     1) if modeD=0
#     C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 <= 1 }
#     1) if modeD=1
#     C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#                                     gamma1||d_j||_1 <= 1 }
#     1) if modeD=2
#     C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#                                     gamma1||d_j||_1 + gamma2 FL(d_j) <= 1 }
#
# Potentially, n can be very large with this algorithm.
#
# Inputs:
#     X: double m x n matrix (input signals)
#         m is the signal size
#         n is the number of signals to decompose
#     D: (optional) double m x p matrix (dictionary)
#         p is the number of elements in the dictionary
#         When D is not provided, the dictionary is initialized
#         with random elements from the training set.
#     K: (size of the dictionary, optional is D is provided)
#     lambda1: (parameter)
#     iter: (number of iterations). If a negative number is
#         provided it will perform the computation during the
#         corresponding number of seconds. For instance iter=-5
#         learns the dictionary during 5 seconds.
#     mode: (optional, see above, by default 2)
#     modeD: (optional, see above, by default 0)
#     posD: (optional, adds positivity constraints on the
#         dictionary, false by default, not compatible with
#         modeD=2)
#     gamma1: (optional parameter for modeD >= 1)
#     gamma2: (optional parameter for modeD = 2)
#     batchsize: (optional, size of the minibatch, by default
#         512)
#     iter_updateD: (optional, number of BCD iterations for the dictionary
#         update step, by default 1)
#     modeParam: (optimization mode).
#         1) if modeParam=0, the optimization uses the
#             parameter free strategy of the ICML paper
#         2) if modeParam=1, the optimization uses the
#             parameters rho as in arXiv:0908.0050
#         3) if modeParam=2, the optimization uses exponential
#             decay weights with updates of the form
#             A_{t} <- rho A_{t-1} + alpha_t alpha_t^T
#     rho: (optional) tuning parameter (see paper arXiv:0908.0050)
#     t0: (optional) tuning parameter (see paper arXiv:0908.0050)

```

```

#   clean: (optional, true by default. prunes
#         automatically the dictionary from unused elements).
#   numThreads: (optional, number of threads for exploiting
#               multi-core / multi-cpus. By default, it takes the value -1,
#               which automatically selects all the available CPUs/cores).
#   expand:      undocumented; modify at your own risks!
#   whiten:     undocumented; modify at your own risks!
#   stochastic_deprecated:  undocumented; modify at your own risks!
#   batch:      undocumented; modify at your own risks!
#   log_deprecated:  undocumented; modify at your own risks!
#   logName:    undocumented; modify at your own risks!
#
# Output:
#   D:      double m x p matrix   (dictionary)
#   model:  the model as A B iter
#   D <- spams.trainDL_Memory(X,...)
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting (even though the output alpha is double
#     precision)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
I = readPNG(paste(.imagesDir,'lena.png',sep= '/'))
if (length(dim(I)) == 3) {
  A = matrix(I,nrow = nrow(I),ncol = 3 * ncol(I))
} else {
  A = I
}

m = 8;n = 8;
X = spams.im2col_sliding(A,m,n)

X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
X = X[,seq(from = 1,to = ncol(X),by = 10)]
lambda1 = 0.15

##### FIRST EXPERIMENT #####
tic = proc.time()
D = spams.trainDL_Memory(X,K = 100,lambda1 = lambda1, numThreads = 4,iter = 100)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

```

```

#### SECOND EXPERIMENT ####
tic = proc.time()
D = spams.trainDL(X,K = 100,lambda1 = lambda1, numThreads = 4,iter = 100)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)

.objective(X,D,lambda1)

```

3.3 Function spams.structTrainDL

This function allows to use spams.trainDL with structured regularization functions for the coefficients α . It internally uses the FISTA algorithm.

```

#
# Name: spams.structTrainDL
#
# Usage: spams.structTrainDL(X,return_model= FALSE,model= NULL,D = NULL,graph = NULL,tree =
#       NULL,
#           numThreads = -1,tol = 0.000001,fixed_step = TRUE,ista = FALSE,
#           batchsize = -1,K= -1,lambda1= NULL,lambda2= 10e-10,lambda3 = 0.,
#           iter=-1,t0=1e-5,regul='none',posAlpha=FALSE,posD=FALSE,expand=FALSE,
#           modeD='L2',whiten=FALSE,clean=TRUE,verbose=TRUE,gamma1=0.,gamma2=0.,
#           rho=1.0,iter_updated=NULL,stochastic_deprecated=FALSE,modeParam=0,
#           batch=FALSE,log_deprecated=FALSE,logName='')
#
# Description:
#   spams.structTrainDL is an efficient implementation of the
#   dictionary learning technique presented in
#
#   "Online Learning for Matrix Factorization and Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   arXiv:0908.0050
#
#   "Online Dictionary Learning for Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   ICML 2009.
#
#   It addresses the dictionary learning problems
#    $\min_{\{D \in C\}} (1/n) \sum_{i=1}^n 0.5 \|x_i - D\alpha_i\|_2^2 + \lambda \text{psi}(\alpha)$ 
#   where the regularization function psi depends on regul
#   (see spams.proximalFlat for the description of psi,
#   and regul below for allowed values of regul)
#
#   C is a convex set verifying
#   1) if modeD=0
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 <= 1 }
#   2) if modeD=1
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#           gamma1||d_j||_1 <= 1 }
#   3) if modeD=2
#       C={ D in Real^{m x p} s.t. forall j, ||d_j||_2^2 + ...
#           gamma1||d_j||_1 + gamma2 FL(d_j) <= 1 }
#

```

```

# 4) if modeD=3
# C={ D in Real^{m x p} s.t. forall j, (1-gamma1)||d_j||_2^2 + ...
# gamma1||d_j||_1 <= 1 }
#
# Potentially, n can be very large with this algorithm.
#
# Inputs:
# X: double m x n matrix (input signals)
#     m is the signal size
#     n is the number of signals to decompose
# return_model:
#     if true the function will return the model
#     as a named list ('A' = A, 'B' = B, 'iter' = n)
# model: NULL or model (as A,B,iter) to use as initialisation
# D: (optional) double m x p matrix (dictionary)
#     p is the number of elements in the dictionary
#     When D is not provided, the dictionary is initialized
#     with random elements from the training set.
# K: (size of the dictionary, optional is D is provided)
# lambda1: (parameter)
# lambda2: (optional, by default 0)
# lambda3: (optional, regularization parameter, 0 by default)
# iter: (number of iterations). If a negative number is
#     provided it will perform the computation during the
#     corresponding number of seconds. For instance iter=-5
#     learns the dictionary during 5 seconds.
# regul: choice of regularization : one of
#     'l0' 'l1' 'l2' 'linf' 'none' 'elastic-net' 'fused-lasso'
#     'graph' 'graph-ridge' 'graph-l2' 'tree-l0' 'tree-l2' 'tree-linf'
# tree: struct (see documentation of spams.proximalTree);
#     needed for regul of graph kind.
# graph: struct (see documentation of spams.proximalGraph);
#     needed for regul of tree kind.
# posAlpha: (optional, adds positivity constraints on the
#     coefficients, false by default.
# modeD: (optional, see above, by default 0)
# posD: (optional, adds positivity constraints on the
#     dictionary, false by default, not compatible with
#     modeD=2)
# gamma1: (optional parameter for modeD >= 1)
# gamma2: (optional parameter for modeD = 2)
# batchsize: (optional, size of the minibatch, by default
#     512)
# iter_updateD: (optional, number of BCD iterations for the dictionary
#     update step, by default 1)
# modeParam: (optimization mode).
#     1) if modeParam=0, the optimization uses the
#         parameter free strategy of the ICML paper
#     2) if modeParam=1, the optimization uses the
#         parameters rho as in arXiv:0908.0050
#     3) if modeParam=2, the optimization uses exponential
#         decay weights with updates of the form
#         A_{t} <- rho A_{t-1} + alpha_t alpha_t^T
# ista: (optional, use ista instead of fista, false by default).

```

```

#   tol: (optional, tolerance for stopping criterion, which is a relative duality gap
#   fixed_step: (deactivate the line search for L in fista and use K instead)
#   rho: (optional) tuning parameter (see paper arXiv:0908.0050)
#   t0: (optional) tuning parameter (see paper arXiv:0908.0050)
#   clean: (optional, true by default. prunes
#           automatically the dictionary from unused elements).
#   verbose: (optional, true by default, increase verbosity)
#   numThreads: (optional, number of threads for exploiting
#               multi-core / multi-cpus. By default, it takes the value -1,
#               which automatically selects all the available CPUs/cores).
#   expand:    undocumented; modify at your own risks!
#   whiten:    undocumented; modify at your own risks!
#   stochastic_deprecated:    undocumented; modify at your own risks!
#   batch:    undocumented; modify at your own risks!
#   log_deprecated:    undocumented; modify at your own risks!
#   logName:    undocumented; modify at your own risks!
#
# Output:
#   D:    double m x p matrix    (dictionary)
#   model: the model as A B iter
#   D <- spams.structTrainDL(X,return_model = FALSE,...)
#   v <- spams.structTrainDL(X,return_model = TRUE,...)
#   D <- v[[1]]
#   model <- v[[2]]
#   A = model[['A']]
#   B = model[['B']]
#   iter = model[['iter']]
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
I = readPNG(paste(.imagesDir, 'boat.png', sep= '/'))
if (length(dim(I)) == 3) {
  A = matrix(I,nrow = nrow(I),ncol = 3 * ncol(I))
} else {
  A = I
}

m = 8;n = 8;
X = spams.im2col_sliding(A,m,n)

X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)

```

```

lambda1 = 0.05
K = 64
tol = 1e-3
iter = 20
##### FIRST EXPERIMENT #####
regul = 'l1'
.printf("with Fista Regression %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
#
regul = 'l2'
.printf("with Fista Regression %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
#
regul = 'elastic-net'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)

# pause :
# readline("pause ")
##### GRAPH
lambda1 = 0.1
tol = 1e-5
K = 10
eta_g = as.vector(c(1, 1, 1, 1, 1),mode='double')
groups = as(matrix(as.vector(c(0, 0, 0, 1, 0,
  0, 0, 0, 0, 0,
  0, 0, 0, 0, 0,
  0, 0, 0, 0, 0,
  0, 0, 1, 0, 0),mode='logical'),ncol = 5,byrow = T),'CsparseMatrix')

groups_var = as(matrix(as.vector(c(1, 0, 0, 0, 0,
  1, 0, 0, 0, 0,
  1, 0, 0, 0, 0,
  1, 1, 0, 0, 0,
  0, 1, 0, 1, 0,
  0, 1, 0, 1, 0,

```

```

0, 1, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 1, 0, 0),mode='logical'),ncol = 5,byrow = T), 'CsparseMatrix')
graph = list('eta_g' = eta_g, 'groups' = groups, 'groups_var' = groups_var)

regul = 'graph'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul, graph = graph)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
#
regul = 'graph-ridge'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul, graph = graph)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
# pause :
# readline("pause ")
##### TREE

lambda1 = 0.001
tol = 1e-5
treedata = "0 1. [] -> 1 4
1 1. [0 1 2] -> 2 3
4 2. [] -> 5 6
2 1. [3 4]
3 2. [5]
5 2. [6 7]
6 2.5 [8] -> 7
7 2.5 [9]
"

own_variables = as.vector(c(0, 0, 3, 5, 6, 6, 8, 9),mode= 'integer')
N_own_variables = as.vector(c(0,3,2,1,0,2,1,1),mode= 'integer')
eta_g = as.vector(c(1,1,1,2,2,2,2.5,2.5),mode = 'double')
groups = matrix(as.vector(c(0,0,0,0,0,0,0,0,0,
  1,0,0,0,0,0,0,0,0,
  0,1,0,0,0,0,0,0,0,
  0,1,0,0,0,0,0,0,0,
  1,0,0,0,0,0,0,0,0,
  0,0,0,0,1,0,0,0,0,
  0,0,0,0,1,0,0,0,0,
  0,0,0,0,0,0,1,0),mode='logical'),ncol = 8,byrow = T)
groups = as(groups, 'CsparseMatrix')
tree = list('eta_g' = eta_g, 'groups' = groups, 'own_variables' = own_variables,
  'N_own_variables' = N_own_variables)

```



```

#
regul = 'tree-l0'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul, tree = tree)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
#
gstruct = spams.groupStructOfString(treedata)
x = spams.treeOfGroupStruct(gstruct)
tree = x[[2]]
regul = 'tree-l2'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul, tree = tree)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)
#
regul = 'tree-linf'
.printf("with Fista %s\n",regul)
tic = proc.time()
D <- spams.structTrainDL(X,K = K,lambda1 = lambda1, tol = tol,numThreads = 4, batchsize = 400,
  iter = iter,regul = regul, tree = tree)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
.objective(X,D,lambda1)

```

3.4 Function nmf

This function is an example on how to use the function `spams.trainDL` for the problem of non-negative matrix factorization formulated in [17]. Note that `spams.trainDL` can be replaced by `spams.trainDL_Memory` in this function for small or medium datasets.

```

#
# Name: nmf
#
# Usage: spams.nmf(X,return_lasso= FALSE,model= NULL,numThreads = -1,batchsize = -1,K= -1,iter
  =-1,
#         t0=1e-5,clean=TRUE,rho=1.0,modeParam=0,batch=FALSE)
#
# Description:
#   spams.trainDL is an efficient implementation of the
#   non-negative matrix factorization technique presented in
#
#   "Online Learning for Matrix Factorization and Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   arXiv:0908.0050

```

```

#
# "Online Dictionary Learning for Sparse Coding"
# by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
# ICML 2009.
#
# Potentially, n can be very large with this algorithm.
#
# Inputs:
# X: double m x n matrix (input signals)
#     m is the signal size
#     n is the number of signals to decompose
# return_lasso:
#     if true the function will return 2 matrices in a list.
# K: (number of required factors)
# iter: (number of iterations). If a negative number
#     is provided it will perform the computation during the
#     corresponding number of seconds. For instance iter=-5
#     learns the dictionary during 5 seconds.
# batchsize: (optional, size of the minibatch, by default
#     512)
# modeParam: (optimization mode).
#     1) if modeParam=0, the optimization uses the
#     parameter free strategy of the ICML paper
#     2) if modeParam=1, the optimization uses the
#     parameters rho as in arXiv:0908.0050
#     3) if modeParam=2, the optimization uses exponential
#     decay weights with updates of the form
#      $A_{t} \leftarrow \rho A_{t-1} + \alpha_t \alpha_t^T$ 
# rho: (optional) tuning parameter (see paper
# arXiv:0908.0050)
# t0: (optional) tuning parameter (see paper
# arXiv:0908.0050)
# clean: (optional, true by default. prunes automatically
#     the dictionary from unused elements).
# batch: (optional, false by default, use batch learning
#     instead of online learning)
# numThreads: (optional, number of threads for exploiting
#     multi-core / multi-cpus. By default, it takes the value -1,
#     which automatically selects all the available CPUs/cores).
# model: struct (optional) learned model for "retraining" the data.
#
# Output:
# U: double m x p matrix
# V: double p x n matrix (optional)
# model: struct (optional) learned model to be used for
#     "retraining" the data.
# U <- spams.nmf(X,return_lasso = FALSE,...)
# v <- spams.nmf(X,return_lasso = TRUE,...)
# U <- v[[1]]
# V <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)

```

```
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
I = readPNG(paste(.imagesDir,'boat.png',sep= '/'))
if (length(dim(I)) == 3) {
  A = matrix(I,nrow = nrow(I),ncol = 3 * ncol(I))
} else {
  A = I
}
m = 16;n = 16;
X = spams.im2col_sliding(A,m,n)
X = X[,seq(from = 1,to = ncol(X),by = 10)]
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)

##### FIRST EXPERIMENT #####
tic = proc.time()
res <- spams.nmf(X,return_lasso= TRUE,K = 49,numThreads=4,iter = -5)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("time of computation for Dictionary Learning: %f\n",t)
U = res[[1]]
V = res[[2]]
.printf("Evaluating cost function...\n")
R = mean(0.5 * colSums((X - U %*% V) ^ 2))
.printf("objective function: %f\n",R)
```

3.5 Function nmsc

This function is an example on how to use the function `spams.trainDL` for the problem of non-negative sparse coding as defined in [14]. Note that `spams.trainDL` can be replaced by `spams.trainDL_Memory` in this function for small or medium datasets.

```
#
# Name: nmf
#
# Usage: spams.nmsc(X,return_lasso= FALSE,model= NULL,lambda1= NULL,numThreads = -1,batchsize =
#         -1,
#         K= -1,iter=-1,t0=1e-5,clean=TRUE,rho=1.0,modeParam=0,batch=FALSE)
#
# Description:
#   spams.trainDL is an efficient implementation of the
#   non-negative sparse coding technique presented in
#
#   "Online Learning for Matrix Factorization and Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   arXiv:0908.0050
#
#   "Online Dictionary Learning for Sparse Coding"
#   by Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro
#   ICML 2009.
#
#   Potentially, n can be very large with this algorithm.
#
```

```

# Inputs:
# X: double m x n matrix (input signals)
#     m is the signal size
#     n is the number of signals to decompose
# return_lasso:
#     if true the function will return 2 matrices in a list.
# K: (number of required factors)
# lambda1: (parameter)
# iter: (number of iterations). If a negative number
#     is provided it will perform the computation during the
#     corresponding number of seconds. For instance iter=-5
#     learns the dictionary during 5 seconds.
# batchSize: (optional, size of the minibatch, by default
#     512)
# modeParam: (optimization mode).
#     1) if modeParam=0, the optimization uses the
#         parameter free strategy of the ICML paper
#     2) if modeParam=1, the optimization uses the
#         parameters rho as in arXiv:0908.0050
#     3) if modeParam=2, the optimization uses exponential
#         decay weights with updates of the form
#          $A_{t} \leftarrow \rho A_{t-1} + \alpha_t \alpha_t^T$ 
# rho: (optional) tuning parameter (see paper
# arXiv:0908.0050)
# t0: (optional) tuning parameter (see paper
# arXiv:0908.0050)
# clean: (optional, true by default. prunes automatically
#     the dictionary from unused elements).
# batch: (optional, false by default, use batch learning
#     instead of online learning)
# numThreads: (optional, number of threads for exploiting
#     multi-core / multi-cpus. By default, it takes the value -1,
#     which automatically selects all the available CPUs/cores).
# model: struct (optional) learned model for "retraining" the data.
#
# Output:
# U: double m x p matrix
# V: double p x n matrix (optional)
# model: struct (optional) learned model to be used for
#     "retraining" the data.
# U <- spams.nnsc(X,return_lasso = FALSE,...)
# v <- spams.nnsc(X,return_lasso = TRUE,...)
# U <- v[[1]]
# V <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

3.6 Function spams.archetypalAnalysis

The function optimizes the archetypal analysis formulation of [7]. It follows the methodology presented in the following paper [37].

```

#
# Name: spams.archetypalAnalysis
#
# Usage: spams.archetypalAnalysis(X,p = 10,ZO = NULL,returnAB = FALSE,robust=FALSE,epsilon=1e
-3,
#
#           computeXtX=TRUE,stepsFISTA=3,stepsAS=50,randominit=TRUE,
#           numThreads=-1)
#
# Description:
#   documentation to appear soon
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#
# Output:
#   Z: double %
#
# Authors:
# Yuansi Chen and Julien MAIRAL, 2014 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function.

4 Sparse Decomposition Toolbox

This toolbox implements several algorithms for solving signal reconstruction problems. It is mostly adapted for solving a large number of small/medium scale problems, but can be also efficient sometimes with large scale ones.

4.1 Function `spams.omp`

This is a fast implementation of the Orthogonal Matching Pursuit algorithm (or forward selection) [27, 35]. Given a matrix of signals $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ in $\mathbb{R}^{m \times n}$ and a dictionary $\mathbf{D} = [\mathbf{d}^1, \dots, \mathbf{d}^p]$ in $\mathbb{R}^{m \times p}$, the algorithm computes a matrix $\mathbf{A} = [\boldsymbol{\alpha}^1, \dots, \boldsymbol{\alpha}^n]$ in $\mathbb{R}^{p \times n}$, where for each column \mathbf{x} of \mathbf{X} , it returns a coefficient vector $\boldsymbol{\alpha}$ which is an approximate solution of the following NP-hard problem

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \quad \text{s.t.} \quad \|\boldsymbol{\alpha}\|_0 \leq L, \quad (2)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\boldsymbol{\alpha}\|_0 \quad \text{s.t.} \quad \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \leq \varepsilon, \quad (3)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda \|\boldsymbol{\alpha}\|_0. \quad (4)$$

For efficiency reasons, the method first computes the covariance matrix $\mathbf{D}^T \mathbf{D}$, then for each signal, it computes $\mathbf{D}^T \mathbf{x}$ and performs the decomposition with a Cholesky-based algorithm (see [6] for instance).

Note that `spams.omp` can return the “greedy” regularization path if needed (see below):

```

#
# Name: spams.omp
#
# Usage: spams.omp(X,D,L = NULL,eps = NULL,lambda1 = NULL,return_reg_path = FALSE,numThreads =
-1)

```

```

#
# Description:
#   spams.omp is an efficient implementation of the
#   Orthogonal Matching Pursuit algorithm. It is optimized
#   for solving a large number of small or medium-sized
#   decomposition problem (and not for a single large one).
#   It first computes the Gram matrix D'D and then perform
#   a Cholesky-based OMP of the input signals in parallel.
#   X=[x^1,...,x^n] is a matrix of signals, and it returns
#   a matrix A=[alpha^1,...,alpha^n] of coefficients.
#
#   it addresses for all columns x of X,
#       min_{alpha} ||alpha||_0 s.t. ||x-Dalpha||_2^2 <= eps
#       or
#       min_{alpha} ||x-Dalpha||_2^2 s.t. ||alpha||_0 <= L
#       or
#       min_{alpha} 0.5||x-Dalpha||_2^2 + lambda1||alpha||_0
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   D: double m x p matrix (dictionary)
#       p is the number of elements in the dictionary
#       All the columns of D should have unit-norm !
#   return_reg_path:
#       if true the function will return 2 matrices in a list.
#   L: (optional, maximum number of elements in each decomposition,
#       min(m,p) by default)
#   eps: (optional, threshold on the squared l2-norm of the residual,
#       0 by default)
#   lambda1: (optional, penalty parameter, 0 by default)
#   numThreads: (optional, number of threads for exploiting
#   multi-core / multi-cpus. By default, it takes the value -1,
#   which automatically selects all the available CPUs/cores).
#
# Output:
#   A: double sparse p x n matrix (output coefficients)
#   path (optional): double dense p x L matrix (regularization path of the first signal)
#   A <- spams.omp(X,D,L,eps,return_reg_path = FALSE,...)
#   v <- spams.omp(X,D,L,eps,return_reg_path = TRUE,...)
#   A <- v[[1]]
#   path <- v[[2]]
#
# Authors:
#   Julien MAIRAL, 2009 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting (even though the output alpha is double
#     precision)
#   - Passing an int32 vector of length n to L provides

```

```
# a different parameter L for each input signal x_i
# - Passing a double vector of length n to eps and or lambda1
# provides a different parameter eps (or lambda1) for each input signal x_i
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
set.seed(0)
.printf("test omp\n")
X = matrix(rnorm(64 * 100000),nrow = 64,ncol = 100000,byrow = FALSE)
D = matrix(rnorm(64 * 200),nrow = 64,ncol = 200,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
L = 10
eps =0.1
numThreads = -1

tic = proc.time()
alpha = spams.omp(X,D,L=L,eps=eps,return_reg_path = FALSE,numThreads = numThreads)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)

#####
# Regularization path of a single signal
#####
X = matrix(rnorm(64 * 1),nrow = 64,ncol = 1,byrow = FALSE)
D = matrix(rnorm(64 * 10),nrow = 64,ncol = 10,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
L = 5
res = spams.omp(X,D,L = L,eps = eps,return_reg_path = TRUE,numThreads = numThreads)
alpha = res[[1]]
path = res[[2]]
```

4.2 Function spams.ompMask

This is a variant of spams.omp with the possibility of handling a binary mask. Given a binary mask $\mathbf{B} = [\beta^1, \dots, \beta^n]$ in $\{0, 1\}^{m \times n}$, it returns a matrix $\mathbf{A} = [\alpha^1, \dots, \alpha^n]$ such that for every column \mathbf{x} of \mathbf{X} , β of \mathbf{B} , it computes a column α of \mathbf{A} by addressing

$$\min_{\alpha \in \mathbb{R}^p} \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 \quad \text{s.t.} \quad \|\alpha\|_0 \leq L, \quad (5)$$

or

$$\min_{\alpha \in \mathbb{R}^p} \|\alpha\|_0 \quad \text{s.t.} \quad \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 \leq \varepsilon \frac{\|\beta\|_0}{m}, \quad (6)$$

or

$$\min_{\alpha \in \mathbb{R}^p} \frac{1}{2} \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 + \lambda \|\alpha\|_0. \quad (7)$$

where $\text{diag}(\beta)$ is a diagonal matrix with the entries of β on the diagonal.

```
#
# Name: spams.ompMask
#
# Usage: spams.ompMask(X,D,B,L = NULL,eps = NULL,lambda1 = NULL,return_reg_path = FALSE,
# numThreads = -1)
```

```

#
# Description:
#   spams.ompMask is a variant of spams.omp that allow using
#   a binary mask B
#
#   for all columns x of X, and columns beta of B, it computes a column
#   alpha of A by addressing
#   min_{alpha} ||alpha||_0 s.t. ||diag(beta)*(x-Dalpha)||_2^2
#                                     <= eps*||beta||_0/m
#
#   or
#   min_{alpha} ||diag(beta)*(x-Dalpha)||_2^2 s.t. ||alpha||_0 <= L
#   or
#   min_{alpha} 0.5||diag(beta)*(x-Dalpha)||_2^2 + lambda1||alpha||_0
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   D: double m x p matrix (dictionary)
#       p is the number of elements in the dictionary
#       All the columns of D should have unit-norm !
#   B: boolean m x n matrix (mask)
#       p is the number of elements in the dictionary
#   return_reg_path:
#       if true the function will return 2 matrices in a list.
#   L: (optional, maximum number of elements in each decomposition,
#       min(m,p) by default)
#   eps: (optional, threshold on the squared l2-norm of the residual,
#       0 by default)
#   lambda1: (optional, penalty parameter, 0 by default)
#   numThreads: (optional, number of threads for exploiting
#   multi-core / multi-cpus. By default, it takes the value -1,
#   which automatically selects all the available CPUs/cores).
#
# Output:
#   A: double sparse p x n matrix (output coefficients)
#   path (optional): double dense p x L matrix
#                       (regularization path of the first signal)
#   A <- spams.ompMask(X,D,B,L,eps,return_reg_path = FALSE
#   ,...)
#   v <- spams.ompMask(X,D,B,L,eps,return_reg_path = TRUE
#   ,...)
#
#   A <- v[[1]]
#   path <- v[[2]]
#
# Authors:
#   Julien MAIRAL, 2010 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting (even though the output alpha is double
#   precision)

```



```
# - Passing an int32 vector of length n to L provides
# a different parameter L for each input signal x_i
# - Passing a double vector of length n to eps and or lambda1
# provides a different parameter eps (or lambda1) for each input signal x_i
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
set.seed(0)
.printf("test ompMask\n")
X = matrix(rnorm(100 * 100),nrow = 100,ncol = 100,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
D = matrix(rnorm(100 * 20),nrow = 100,ncol = 20,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
mask = (X > 0) # generating a binary mask
L = 20
eps =0.01
numThreads = -1

tic = proc.time()
alpha = spams.ompMask(X,D,mask,L = L,eps = eps,return_reg_path = FALSE,numThreads = numThreads)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)
```

4.3 Function mexRidgeRegression

This is a ridge regression solver using a conjugate gradient solver.

```
#
# The R function is not yet implemented.
#
```

The following piece of code illustrates how to use this function.

4.4 Function spams.lasso

This is a fast implementation of the LARS algorithm [9] (variant for solving the Lasso) for solving the Lasso or Elastic-Net. Given a matrix of signals $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ in $\mathbb{R}^{m \times n}$ and a dictionary \mathbf{D} in $\mathbb{R}^{m \times p}$, depending on the input parameters, the algorithm returns a matrix of coefficients $\mathbf{A} = [\boldsymbol{\alpha}^1, \dots, \boldsymbol{\alpha}^n]$ in $\mathbb{R}^{p \times n}$ such that for every column \mathbf{x} of \mathbf{X} , the corresponding column $\boldsymbol{\alpha}$ of \mathbf{A} is the solution of

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \quad \text{s.t.} \quad \|\boldsymbol{\alpha}\|_1 \leq \lambda, \quad (8)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\boldsymbol{\alpha}\|_1 \quad \text{s.t.} \quad \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \leq \lambda, \quad (9)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda \|\boldsymbol{\alpha}\|_1 + \frac{\lambda_2}{2} \|\boldsymbol{\alpha}\|_2^2. \quad (10)$$

For efficiency reasons, the method first compute the covariance matrix $\mathbf{D}^T \mathbf{D}$, then for each signal, it computes $\mathbf{D}^T \mathbf{x}$ and performs the decomposition with a Cholesky-based algorithm (see [9] for instance). The implementation has also an option to add **positivity constraints** on the solutions $\boldsymbol{\alpha}$. When the solution is very sparse and the problem size is reasonable, this approach can be very efficient. Moreover, it gives the solution with

an exact precision, and its performance does not depend on the correlation of the dictionary elements, except when the solution is not unique (the algorithm breaks in this case).

Note that `spams.lasso` can return the whole regularization path of the first signal \mathbf{x}_1 and can handle implicitly the matrix \mathbf{D} if the quantities $\mathbf{D}^T\mathbf{D}$ and $\mathbf{D}^T\mathbf{x}$ are passed as an argument, see below:

```
#
# Name: spams.lasso
#
# Usage: spams.lasso(X,D= NULL,Q = NULL,q = NULL,return_reg_path = FALSE,L= -1,lambda1= NULL,
#         lambda2= 0.,mode= 'PENALTY',pos= FALSE,ols= FALSE,numThreads= -1,
#         max_length_path= -1,verbose=FALSE,cholesky= FALSE)
#
# Description:
#   spams.lasso is an efficient implementation of the
#   homotopy-LARS algorithm for solving the Lasso.
#
#   If the function is called this way spams.lasso(X,D = D, Q = NULL,...),
#   it aims at addressing the following problems
#   for all columns x of X, it computes one column alpha of A
#   that solves
#     1) when mode=0
#       min_{alpha} ||x-Dalpha||_2^2 s.t. ||alpha||_1 <= lambda1
#     2) when mode=1
#       min_{alpha} ||alpha||_1 s.t. ||x-Dalpha||_2^2 <= lambda1
#     3) when mode=2
#       min_{alpha} 0.5||x-Dalpha||_2^2 + lambda1||alpha||_1 +0.5 lambda2||alpha||_2^2
#
#   If the function is called this way spams.lasso(X,D = NULL, Q = Q, q = q,...),
#   it solves the above optimisation problem, when Q=D'D and q=D'x.
#
#   Possibly, when pos=true, it solves the previous problems
#   with positivity constraints on the vectors alpha
#
# Inputs:
#   X: double m x n matrix   (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   D: double m x p matrix   (dictionary)
#       p is the number of elements in the dictionary
#   Q:           p x p double matrix (Q = D'D)
#   q:           p x n double matrix (q = D'X)
#   verbose:     verbose mode
#   return_reg_path:
#       if true the function will return 2 matrices in a list.
#   lambda1: (parameter)
#   lambda2: (optional parameter for solving the Elastic-Net)
#             for mode=0 and mode=1, it adds a ridge on the Gram Matrix
#   L: (optional), maximum number of steps of the homotopy algorithm (can
#       be used as a stopping criterion)
#   pos: (optional, adds non-negativity constraints on the
#       coefficients, false by default)
#   mode: (see above, by default: 2)
#   numThreads: (optional, number of threads for exploiting
#       multi-core / multi-cpus. By default, it takes the value -1,
#       which automatically selects all the available CPUs/cores).
```

```

#   cholesky: (optional, default false), choose between Cholesky
#   implementation or one based on the matrix inversion Lemma
#   ols: (optional, default false), perform an orthogonal projection
#   before returning the solution.
#   max_length_path: (optional) maximum length of the path, by default 4*p
#
# Output:
#   A: double sparse p x n matrix (output coefficients)
#   path: optional, returns the regularisation path for the first signal
#   A <- spams.lasso(X,return_reg_path = FALSE,...)
#   v <- spams.lasso(X,return_reg_path = TRUE,...)
#   A <- v[[1]]
#   path <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting (even though the output alpha is double
#     precision)
#
# Examples:
#   m <- 5;n <- 10;nD <- 5
#   set.seed(0)
#   X = matrix(rnorm(m * n),nrow = m,ncol = n,byrow = FALSE)
#   X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=TRUE)
#   D = matrix(rnorm(m * nD),nrow = m,ncol = nD,byrow = FALSE)
#   D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=TRUE)
#   alpha = spams.lasso(X,D = D,return_reg_path = FALSE,lambda1 = 0.15)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
.printf("test lasso\n")
#####
# Decomposition of a large number of signals
#####
# data generation
X = matrix(rnorm(100 * 100000),nrow = 100,ncol = 100000,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
D = matrix(rnorm(100 * 200),nrow = 100,ncol = 200,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
tic = proc.time()
alpha = spams.lasso(X,D,return_reg_path = FALSE,lambda1 = 0.15,numThreads = -1,mode = 'PENALTY'
)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)
#####

```

```

# Regularization path of a single signal
#####
X = matrix(rnorm(64),nrow = 64,ncol = 1,byrow = FALSE)
D = matrix(rnorm(64 * 10),nrow = 64,ncol = 10,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
res = spams.lasso(X,D,return_reg_path = TRUE,lambda1 = 0.15,numThreads = -1,mode = 'PENALTY' )
alpha = res[[1]]
path = res[[2]]

```

4.5 Function spams.lassoWeighted

This is a fast implementation of a weighted version of LARS [9]. Given a matrix of signals $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ in $\mathbb{R}^{m \times n}$, a matrix of weights $\mathbf{W} = [\mathbf{w}^1, \dots, \mathbf{w}^n] \in \mathbb{R}^{p \times n}$, and a dictionary \mathbf{D} in $\mathbb{R}^{m \times p}$, depending on the input parameters, the algorithm returns a matrix of coefficients $\mathbf{A} = [\boldsymbol{\alpha}^1, \dots, \boldsymbol{\alpha}^n]$ in $\mathbb{R}^{p \times n}$, such that for every column \mathbf{x} of \mathbf{X} , \mathbf{w} of \mathbf{W} , it computes a column $\boldsymbol{\alpha}$ of \mathbf{A} , which is the solution of

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \quad \text{s.t.} \quad \|\text{diag}(\mathbf{w})\boldsymbol{\alpha}\|_1 \leq \lambda, \quad (11)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \|\text{diag}(\mathbf{w})\boldsymbol{\alpha}\|_1 \quad \text{s.t.} \quad \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 \leq \lambda, \quad (12)$$

or

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda \|\text{diag}(\mathbf{w})\boldsymbol{\alpha}\|_1. \quad (13)$$

The implementation has also an option to add **positivity constraints** on the solutions $\boldsymbol{\alpha}$. This function is potentially useful for implementing efficiently the randomized Lasso of [28], or reweighted- ℓ_1 schemes [4].

```

#
# Name: spams.lassoWeighted.
#
# Usage: spams.lassoWeighted(X,D,W,L= -1,lambda1= NULL,mode= 'PENALTY',pos= FALSE,numThreads=
#       -1,
#               verbose=FALSE)
#
# Description:
#   spams.lassoWeighted is an efficient implementation of the
#   LARS algorithm for solving the weighted Lasso. It is optimized
#   for solving a large number of small or medium-sized
#   decomposition problem (and not for a single large one).
#   It first computes the Gram matrix D'D and then perform
#   a Cholesky-based OMP of the input signals in parallel.
#   For all columns x of X, and w of W, it computes one column alpha of A
#   which is the solution of
#   1) when mode=0
#       min_{alpha} ||x-Dalpha||_2^2   s.t.
#                                   ||diag(w)alpha||_1 <= lambda1
#   2) when mode=1
#       min_{alpha} ||diag(w)alpha||_1   s.t.
#                                   ||x-Dalpha||_2^2 <= lambda1
#   3) when mode=2
#       min_{alpha} 0.5||x-Dalpha||_2^2 +
#                                   lambda1||diag(w)alpha||_1
#   Possibly, when pos=true, it solves the previous problems
#   with positivity constraints on the vectors alpha
#
# Inputs:

```

```

# X: double m x n matrix (input signals)
#     m is the signal size
#     n is the number of signals to decompose
# D: double m x p matrix (dictionary)
#     p is the number of elements in the dictionary
# W: double p x n matrix (weights)
# verbose: verbose mode
# lambda1: (parameter)
# L: (optional, maximum number of elements of each
# decomposition)
# pos: (optional, adds positivity constraints on the
# coefficients, false by default)
# mode: (see above, by default: 2)
# numThreads: (optional, number of threads for exploiting
# multi-core / multi-cpus. By default, it takes the value -1,
# which automatically selects all the available CPUs/cores).
#
# Output:
# A: double sparse p x n matrix (output coefficients)
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
# this function admits a few experimental usages, which have not
# been extensively tested:
# - single precision setting (even though the output alpha is double
# precision)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
.printf("test lasso weighted\n")
#####
# Decomposition of a large number of signals
#####
# data generation
X = matrix(rnorm(64 * 10000),nrow = 64,ncol = 10000,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
D = matrix(rnorm(64 * 256),nrow = 64,ncol = 256,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
W = matrix(runif(ncol(D) * ncol(X),0,1),nrow = ncol(D),ncol = ncol(X),byrow = FALSE)
tic = proc.time()
alpha = spams.lassoWeighted(X,D,W,lambda1 = 0.15,numThreads = -1,mode = 'PENALTY')
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)

```

4.6 Function spams.lassoMask

This is a variant of spams.lasso with the possibility of adding a mask $\mathbf{B} = [\beta^1, \dots, \beta^m]$, as in spams.ompMask. For every column \mathbf{x} of \mathbf{X} , β of \mathbf{B} , it computes a column α of \mathbf{A} , which is the solution of

$$\min_{\alpha \in \mathbb{R}^p} \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 \quad \text{s.t.} \quad \|\alpha\|_1 \leq \lambda, \quad (14)$$

or

$$\min_{\alpha \in \mathbb{R}^p} \|\alpha\|_1 \quad \text{s.t.} \quad \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 \leq \lambda \frac{\|\beta\|_0}{m}, \quad (15)$$

or

$$\min_{\alpha \in \mathbb{R}^p} \frac{1}{2} \|\text{diag}(\beta)(\mathbf{x} - \mathbf{D}\alpha)\|_2^2 + \lambda \frac{\|\beta\|_0}{m} \|\alpha\|_1 + \frac{\lambda_2}{2} \|\alpha\|_2^2. \quad (16)$$

```
#
# Name: spams.lassoMask
#
# Usage: spams.lassoMask(X,D,B,L= -1,lambda1= NULL,lambda2= 0.,mode= 'PENALTY',pos= FALSE,
#         numThreads= -1,verbose=FALSE)
#
# Description:
#   spams.lasso is a variant of spams.lasso that handles
#   binary masks. It aims at addressing the following problems
#   for all columns x of X, and beta of B, it computes one column alpha of A
#   that solves
#     1) when mode=0
#         min_{alpha} ||diag(beta)(x-Dalpha)||_2^2 s.t. ||alpha||_1 <= lambda1
#     2) when mode=1
#         min_{alpha} ||alpha||_1 s.t. ||diag(beta)(x-Dalpha)||_2^2
#                                     <= lambda1*||beta||_0/m
#     3) when mode=2
#         min_{alpha} 0.5||diag(beta)(x-Dalpha)||_2^2 +
#                                     lambda1*(||beta||_0/m)*||alpha||_1 +
#                                     (lambda2/2)||alpha||_2^2
#   Possibly, when pos=true, it solves the previous problems
#   with positivity constraints on the vectors alpha
#
# Inputs:
#   X: double m x n matrix   (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   D: double m x p matrix   (dictionary)
#       p is the number of elements in the dictionary
#   B: boolean m x n matrix   (mask)
#       p is the number of elements in the dictionary
#   verbose:      verbose mode
#   lambda1:      (parameter)
#   L: (optional, maximum number of elements of each
#       decomposition)
#   pos: (optional, adds positivity constraints on the
#       coefficients, false by default)
#   mode: (see above, by default: 2)
#   lambda2:      (optional parameter for solving the Elastic-Net)
#                 for mode=0 and mode=1, it adds a ridge on the Gram Matrix
#   numThreads:  (optional, number of threads for exploiting
#                 multi-core / multi-cpus. By default, it takes the value -1,
#                 which automatically selects all the available CPUs/cores).
```

```

#
# Output:
#     A: double sparse p x n matrix (output coefficients)
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#     this function admits a few experimental usages, which have not
#     been extensively tested:
#     - single precision setting (even though the output alpha is double
#       precision)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
.printf("test lassoMask\n")
#####
# Decomposition of a large number of signals
#####
# data generation
X = matrix(rnorm(100 * 100),nrow = 100,ncol = 100,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
D = matrix(rnorm(100 * 20),nrow = 100,ncol = 20,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
mask = (X > 0) # generating a binary mask
tic = proc.time()
alpha = spams.lassoMask(X,D,mask,lambda1 = 0.15,numThreads = -1,mode = 'PENALTY')
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)

```

4.7 Function spams.cd

Coordinate-descent approach for solving Eq. (10) and Eq. (9). Note that unlike `spams.lasso`, it is not implemented to solve the Elastic-Net formulation. To solve Eq. (9), the algorithm solves a sequence of problems of the form (10) using simple heuristics. Coordinate descent is very simple and in practice very powerful. It performs better when the correlation between the dictionary elements is small.

```

#
# Name: spams.cd
#
# Usage: spams.cd(X,D,A0,lambda1 = NULL,mode= 'PENALTY',itermax=100,tol = 0.001,numThreads =-1)
#
# Description:
#     spams.cd addresses l1-decomposition problem with a
#     coordinate descent type of approach.
#     It is optimized for solving a large number of small or medium-sized
#     decomposition problem (and not for a single large one).
#     It first computes the Gram matrix D'D.
#     This method is particularly well adapted when there is low
#     correlation between the dictionary elements and when one can benefit

```

```

# from a warm restart.
# It aims at addressing the two following problems
# for all columns x of X, it computes a column alpha of A such that
# 2) when mode=1
#   min_{alpha} ||alpha||_1 s.t. ||x-Dalpha||_2^2 <= lambda1
#   For this constraint setting, the method solves a sequence of
#   penalized problems (corresponding to mode=2) and looks
#   for the corresponding Lagrange multiplier with a simple but
#   efficient heuristic.
# 3) when mode=2
#   min_{alpha} 0.5||x-Dalpha||_2^2 + lambda1||alpha||_1
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#   D: double m x p matrix (dictionary)
#       p is the number of elements in the dictionary
#       All the columns of D should have unit-norm !
#   A0: double sparse p x n matrix (initial guess)
#   lambda1: (parameter)
#   mode: (optional, see above, by default 2)
#   itermax: (maximum number of iterations)
#   tol: (tolerance parameter)
#   numThreads: (optional, number of threads for exploiting
#   multi-core / multi-cpus. By default, it takes the value -1,
#   which automatically selects all the available CPUs/cores).
#
# Output:
#   A: double sparse p x n matrix (output coefficients)
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting (even though the output alpha
#     is double precision)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
X = matrix(rnorm(64 * 100),nrow = 64,ncol = 100,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X),ncol(X),byrow=T)
D = matrix(rnorm(64 * 100),nrow = 64,ncol = 100,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D),ncol(D),byrow=T)
# parameter of the optimization procedure are chosen
lambda1 = 0.015
mode = 'PENALTY'
tic = proc.time()

```



```

alpha = spams.lasso(X,D,lambda1 = lambda1,mode = mode,numThreads = 4)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second for LARS\n", (ncol(X) / t))
E = mean(0.5 * colSums((X - D %*% alpha) ^ 2) + lambda1 * colSums(abs(alpha)))
.printf("Objective function for LARS: %g\n",E)
tol = 0.001
itermax = 1000
A0 = as(matrix(c(0),nrow = nrow(alpha),ncol = ncol(alpha)), 'CsparseMatrix')
tic = proc.time()
alpha2 = spams.cd(X,D,A0,lambda1 = lambda1,mode = mode,tol = tol, itermax = itermax,numThreads
= 4)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second for CD\n", (ncol(X) / t))
E = mean(0.5 * colSums((X - D %*% alpha2) ^ 2) + lambda1 * colSums(abs(alpha2)))
.printf("Objective function for CD: %g\n",E)
.printf("With Random Design, CD can be much faster than LARS\n")

```

4.8 Function spams.somp

This is a fast implementation of the Simultaneous Orthogonal Matching Pursuit algorithm. Given a set of matrices $\mathbf{X} = [\mathbf{X}^1, \dots, \mathbf{X}^n]$ in $\mathbb{R}^{m \times N}$, where the \mathbf{X}^i 's are in $\mathbb{R}^{m \times n_i}$, and a dictionary \mathbf{D} in $\mathbb{R}^{m \times p}$, the algorithm returns a matrix of coefficients $\mathbf{A} = [\mathbf{A}^1, \dots, \mathbf{A}^n]$ in $\mathbb{R}^{p \times N}$ which is an approximate solution of the following NP-hard problem

$$\forall i \min_{\mathbf{A}^i \in \mathbb{R}^{p \times n_i}} \|\mathbf{X}^i - \mathbf{D}\mathbf{A}^i\|_F^2 \text{ s.t. } \|\mathbf{A}^i\|_{0,\infty} \leq L. \quad (17)$$

or

$$\forall i \min_{\mathbf{A}^i \in \mathbb{R}^{p \times n_i}} \|\mathbf{A}^i\|_{0,\infty} \text{ s.t. } \|\mathbf{X}^i - \mathbf{D}\mathbf{A}^i\|_F^2 \leq \epsilon n_i. \quad (18)$$

To be efficient, the method first compute the covariance matrix $\mathbf{D}^T \mathbf{D}$, then for each signal, it computes $\mathbf{D}^T \mathbf{X}^i$ and performs the decomposition with a Cholesky-based algorithm.

```

#
# Name: spams.somp
# (this function has not been intensively tested).
#
# Usage: spams.somp(X,D,list_groups,L = NULL,eps = 0.,numThreads = -1)
#
# Description:
#   spams.somp is an efficient implementation of a
#   Simultaneous Orthogonal Matching Pursuit algorithm. It is optimized
#   for solving a large number of small or medium-sized
#   decomposition problem (and not for a single large one).
#   It first computes the Gram matrix D'D and then perform
#   a Cholesky-based OMP of the input signals in parallel.
#   It aims at addressing the following NP-hard problem
#
#   X is a matrix structured in groups of signals, which we denote
#   by X=[X_1,...,X_n]
#
#   for all matrices X_i of X,
#       min_{A_i} ||A_i||_{0,infty} s.t. ||X_i-D A_i||_2^2 <= eps*n_i
#       where n_i is the number of columns of X_i
#

```

```

#      or
#
#      min_{A_i} ||X_i - D A_i||_2^2 s.t. ||A_i||_{0,infty} <= L
#
# Inputs:
#      X: double m x N matrix (input signals)
#          m is the signal size
#          N is the total number of signals
#      D: double m x p matrix (dictionary)
#          p is the number of elements in the dictionary
#          All the columns of D should have unit-norm !
#      list_groups : int32 vector containing the indices (starting at 0)
#                   of the first elements of each groups.
#      L: (maximum number of elements in each decomposition)
#      eps: (threshold on the squared l2-norm of the residual
#      numThreads: (optional, number of threads for exploiting
#                   multi-core / multi-cpus. By default, it takes the value -1,
#                   which automatically selects all the available CPUs/cores).
#
# Output:
#      alpha: double sparse p x N matrix (output coefficients)
#
# Authors:
#      Julien MAIRAL, 2010 (spams, matlab interface and documentation)
#      Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#      this function admits a few experimental usages, which have not
#      been extensively tested:
#      - single precision setting (even though the output alpha is double
#        precision)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
X = matrix(rnorm(64 * 10000),nrow = 64,ncol = 10000,byrow = FALSE)
D = matrix(rnorm(64 * 200),nrow = 64,ncol = 200,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
ind_groups = as.vector(seq(from = 0,to = 9999,by = 10),mode= 'integer')
tic = proc.time()
alpha = spams.somp(X,D,ind_groups,L = 10,eps = 0.1,numThreads=-1)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)

```

4.9 Function spams.l1L2BCD

This is a fast implementation of a simultaneous signal decomposition formulation. Given a set of matrices $\mathbf{X} = [\mathbf{X}^1, \dots, \mathbf{X}^n]$ in $\mathbb{R}^{m \times N}$, where the \mathbf{X}^i 's are in $\mathbb{R}^{m \times n_i}$, and a dictionary \mathbf{D} in $\mathbb{R}^{m \times p}$, the algorithm returns a matrix of coefficients $\mathbf{A} = [\mathbf{A}^1, \dots, \mathbf{A}^n]$ in $\mathbb{R}^{p \times N}$ which is an approximate solution of the following NP-hard problem

$$\forall i \quad \min_{\mathbf{A}^i \in \mathbb{R}^{p \times n_i}} \|\mathbf{X}^i - \mathbf{D}\mathbf{A}^i\|_F^2 \quad \text{s.t.} \quad \|\mathbf{A}^i\|_{1,2} \leq \frac{\lambda}{n_i}. \quad (19)$$

or

$$\forall i \min_{\mathbf{A}^i \in \mathbb{R}^{p \times n_i}} \|\mathbf{A}^i\|_{1,2} \text{ s.t. } \|\mathbf{X}^i - \mathbf{D}\mathbf{A}^i\|_F^2 \leq \lambda n_i. \quad (20)$$

To be efficient, the method first compute the covariance matrix $\mathbf{D}^T\mathbf{D}$, then for each signal, it computes $\mathbf{D}^T\mathbf{X}^i$ and performs the decomposition with a Cholesky-based algorithm.

```
#
# Name: spams.l1L2BCD
# (this function has not been intensively tested).
#
# Usage: spams.l1L2BCD(X,D,alpha0,list_groups,lambda1 = NULL,mode= 'PENALTY',itermax = 100,
#         tol = 1e-3,numThreads = -1)
#
# Description:
#   spams.l1L2BCD is a solver for a
#   Simultaneous signal decomposition formulation based on block
#   coordinate descent.
#
#   X is a matrix structured in groups of signals, which we denote
#   by X=[X_1,...,X_n]
#
#   if mode=2, it solves
#       for all matrices X_i of X,
#       min_{A_i} 0.5||X_i-D A_i||_2^2 + lambda1/sqrt(n_i)||A_i||_{1,2}
#       where n_i is the number of columns of X_i
#   if mode=1, it solves
#       min_{A_i} ||A_i||_{1,2} s.t. ||X_i-D A_i||_2^2 <= n_i lambda1
#
# Inputs:
#   X: double m x N matrix   (input signals)
#       m is the signal size
#       N is the total number of signals
#   D: double m x p matrix   (dictionary)
#       p is the number of elements in the dictionary
#   alpha0: double dense p x N matrix (initial solution)
#   list_groups : int32 vector containing the indices (starting at 0)
#       of the first elements of each groups.
#   lambda1: (regularization parameter)
#   mode: (see above, by default 2)
#   itermax: (maximum number of iterations, by default 100)
#   tol: (tolerance parameter, by default 0.001)
#   numThreads: (optional, number of threads for exploiting
#       multi-core / multi-cpus. By default, it takes the value -1,
#       which automatically selects all the available CPUs/cores).
#
# Output:
#   alpha: double sparse p x N matrix (output coefficients)
#
# Authors:
#   Julien MAIRAL, 2010 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
```

```
# - single precision setting (even though the output alpha is double
# precision)
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
set.seed(0)
X = matrix(rnorm(64 * 100),nrow = 64,ncol = 100,byrow = FALSE)
D = matrix(rnorm(64 * 200),nrow = 64,ncol = 200,byrow = FALSE)
D = D / matrix(rep(sqrt(colSums(D*D)),nrow(D)),nrow(D),ncol(D),byrow=T)
ind_groups = as.vector(seq(from = 0,to = ncol(X) - 1,by = 10),mode= 'integer')#indices of the
  first signals in each group
itermax = 100
tol = 1e-3
mode = 'PENALTY'
lambda1 = 0.15 # squared norm of the residual should be less than 0.1
numThreads = -1 # number of processors/cores to use the default choice is -1
  # and uses all the cores of the machine

alpha0 = matrix(c(0),nrow = ncol(D), ncol = ncol(X),byrow = FALSE)
tic = proc.time()
alpha = spams.l1L2BCD(X,D,alpha0,ind_groups,lambda1 = lambda1,mode = mode,itermax = itermax,tol
  = tol,numThreads = numThreads)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf("%f signals processed per second\n",as.double(ncol(X)) / t)
```

4.10 Function spams.sparseProject

This is a multi-purpose function, implementing fast algorithms for projecting on convex sets, but it also solves the fused lasso signal approximation problem. The proposed method is detailed in [21]. The main problems addressed by this function are the following: Given a matrix $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n]$ in $\mathbb{R}^{m \times n}$, it finds a matrix $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n]$ in $\mathbb{R}^{m \times n}$ so that for all column \mathbf{u} of \mathbf{U} , it computes a column \mathbf{v} of \mathbf{V} solving

$$\min_{\mathbf{v} \in \mathbb{R}^m} \|\mathbf{u} - \mathbf{v}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{v}\|_1 \leq \tau, \quad (21)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^m} \|\mathbf{u} - \mathbf{v}\|_2^2 \quad \text{s.t.} \quad \lambda_1 \|\mathbf{v}\|_1 + \lambda_2 \|\mathbf{v}\|_2^2 \leq \tau, \quad (22)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^m} \|\mathbf{u} - \mathbf{v}\|_2^2 \quad \text{s.t.} \quad \lambda_1 \|\mathbf{v}\|_1 + \lambda_2 \|\mathbf{v}\|_2^2 + \lambda_3 FL(\mathbf{v}) \leq \tau, \quad (23)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda_1 \|\mathbf{v}\|_1 + \lambda_2 \|\mathbf{v}\|_2^2 + \lambda_3 FL(\mathbf{v}). \quad (24)$$

Note that for the two last cases, the method performs a small approximation. The method follows the regularization path, goes from one kink to another, and stop whenever the constraint is not satisfied anymore. The solution returned by the algorithm is the one obtained at the last kink of the regularization path, which is in practice close, but not exactly the same as the solution. This will be corrected in a future release of the toolbox.

```
#
# Name: spams.sparseProject
#
# Usage: spams.sparseProject(U,thrs = 1.0,mode = 1,lambda1 = 0.0,lambda2 = 0.0,lambda3 = 0.0,
```

```

#             pos = FALSE,numThreads = -1)
#
# Description:
#   spams.sparseProject solves various optimization
#   problems, including projections on a few convex sets.
#   It aims at addressing the following problems
#   for all columns u of U in parallel
#   1) when mode=1 (projection on the l1-ball)
#       min_v ||u-v||_2^2 s.t. ||v||_1 <= thrs
#   2) when mode=2
#       min_v ||u-v||_2^2 s.t. ||v||_2^2 + lamuda1||v||_1 <= thrs
#   3) when mode=3
#       min_v ||u-v||_2^2 s.t. ||v||_1 + 0.5lamuda1||v||_2^2 <= thrs
#   4) when mode=4
#       min_v 0.5||u-v||_2^2 + lamuda1||v||_1 s.t. ||v||_2^2 <= thrs
#   5) when mode=5
#       min_v 0.5||u-v||_2^2 + lamuda1||v||_1 +lamuda2 FL(v) + ...
#                                     0.5lamuda_3 ||v||_2^2
#       where FL denotes a "fused lasso" regularization term.
#   6) when mode=6
#       min_v ||u-v||_2^2 s.t lamuda1||v||_1 +lamuda2 FL(v) + ...
#                                     0.5lamuda3||v||_2^2 <= thrs
#
#       When pos=true and mode <= 4,
#       it solves the previous problems with positivity constraints
#
# Inputs:
#   U: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to project
#   thrs: (parameter)
#   lambda1: (parameter)
#   lambda2: (parameter)
#   lambda3: (parameter)
#   mode: (see above)
#   pos: (optional, false by default)
#   numThreads: (optional, number of threads for exploiting
#       multi-core / multi-cpus. By default, it takes the value -1,
#       which automatically selects all the available CPUs/cores).
#
# Output:
#   V: double m x n matrix (output matrix)
#
# Authors:
#   Julien MAIRAL, 2009 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   this function admits a few experimental usages, which have not
#   been extensively tested:
#   - single precision setting
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
X = matrix(rnorm(20000 * 100),nrow = 20000,ncol = 100,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
.printf( "\n Projection on the l1 ball\n")
tic = proc.time()
X1 = spams.sparseProject(X,numThreads = -1, pos = FALSE,mode= 1,thrs = 2)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf(" Time : %f\n", t)
if (t != 0)
  .printf("%f signals of size %d projected per second\n", (ncol(X) / t),nrow(X))
s = colSums(abs(X1))
.printf("Checking constraint: %f, %f\n",min(s),max(s))

.printf("\n Projection on the Elastic-Net\n")
lambda1 = 0.15
tic = proc.time()
X1 = spams.sparseProject(X,numThreads = -1, pos = FALSE,mode= 2,thrs = 2,lambda1= lambda1)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf(" Time : %f\n", t)
if (t != 0)
  .printf("%f signals of size %d projected per second\n", (ncol(X) / t),nrow(X))
constraints = colSums(X1*X1) + lambda1 * colSums(abs(X1))
.printf("Checking constraint: %f, %f (Projection is approximate : stops at a kink)\n",min(
  constraints),max(constraints))

.printf("\n Projection on the FLSA\n")
lambda1 = 0.7
lambda2 = 0.7
lambda3 = 0.7
X = matrix(runif(2000 * 100,0,1),nrow = 2000,ncol = 100,byrow = FALSE)
X = X / matrix(rep(sqrt(colSums(X*X)),nrow(X)),nrow(X),ncol(X),byrow=T)
tic = proc.time()
X1 = spams.sparseProject(X,numThreads = -1, pos = FALSE,mode= 6,thrs = 2,lambda1= lambda1,
  lambda2= lambda2,lambda3= lambda3)
tac = proc.time()
t = (tac - tic)[['elapsed']]
.printf(" Time : %f\n", t)
if (t != 0)
  .printf("%f signals of size %d projected per second\n",ncol(X) / t,nrow(X))
# constraints = 0.5 * param[['lambda3']] * colSums(X1*X1) + param[['lambda1']] * colSums(abs(
  X1)) + param[['lambda2']] * abs(X1[3:nrow(X1),]) - colSums(X1[2,])
constraints = 0.5 * lambda3 * colSums(X1*X1) + lambda1 * colSums(abs(X1)) + lambda2 * abs(X1[3:
  nrow(X1),]) - colSums(matrix(X1[2,],nrow=1))
.printf("Checking constraint: %f, %f (Projection is approximate : stops at a kink)\n",min(
  constraints),max(constraints))

```

4.11 Function `spams.decompSimplex`

The function implements an active-set algorithm [37] for solving

$$\min_{\alpha \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x} - \mathbf{D}\alpha\|_2^2 \quad \text{s.t.} \quad \alpha \geq 0 \quad \text{and} \quad \sum_{j=1}^p \alpha[j] = 1.$$

```
#
# Name: spams.decompSimplex
#
# Usage: spams.decompSimplex(X,Z,computeXtX = FALSE,numThreads = -1)
#
# Description:
#   documentation to appear soon
#
# Inputs:
#   X: double m x n matrix (input signals)
#       m is the signal size
#       n is the number of signals to decompose
#
# Output:
#   Z: double %
#
# Authors:
#   Yuansi Chen and Julien MAIRAL, 2014 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
```

5 Proximal Toolbox

The previous toolbox we have presented is well adapted for solving a large number of small and medium-scale sparse decomposition problems with the square loss, which is typical from the classical dictionary learning framework. We now present a new software package that is adapted for solving a wide range of possibly large-scale learning problems, with several combinations of losses and regularization terms. The method implements the proximal methods of [1], and includes the proximal solvers for the tree-structured regularization of [15], and the solver of [22] for general structured sparse regularization. The solver for structured sparse regularization norms includes a C++ max-flow implementation of the push-relabel algorithm of [13], with heuristics proposed by [5].

This implementation also provides robust stopping criteria based on *duality gaps*, which are presented in Appendix A. It can handle intercepts (unregularized variables). The general formulation that our software can solve take the form

$$\min_{\mathbf{w} \in \mathbb{R}^p} [g(\mathbf{w}) \triangleq f(\mathbf{w}) + \lambda\psi(\mathbf{w})],$$

where f is a smooth loss function and ψ is a regularization function. When one optimizes a matrix \mathbf{W} in $\mathbb{R}^{p \times r}$ instead of a vector \mathbf{w} in \mathbb{R}^p , we will write

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}} [g(\mathbf{W}) \triangleq f(\mathbf{W}) + \lambda\psi(\mathbf{W})].$$

Note that the software can possibly handle nonnegativity constraints.

We start by presenting the type of regularization implemented in the software

5.1 Regularization Functions

Our software can handle the following regularization functions ψ for vectors \mathbf{w} in \mathbb{R}^p :

- **The Tikhonov regularization:** $\psi(\mathbf{w}) \triangleq \frac{1}{2}\|\mathbf{w}\|_2^2$.
- **The ℓ_1 -norm:** $\psi(\mathbf{w}) \triangleq \|\mathbf{w}\|_1$.
- **The Elastic-Net:** $\psi(\mathbf{w}) \triangleq \|\mathbf{w}\|_1 + \gamma\|\mathbf{w}\|_2^2$.
- **The Fused-Lasso:** $\psi(\mathbf{w}) \triangleq \|\mathbf{w}\|_1 + \gamma\|\mathbf{w}\|_2^2 + \gamma_2 \sum_{i=1}^{p-1} |\mathbf{w}_{i+1} - \mathbf{w}_i|$.
- **The group Lasso:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_2$, where \mathcal{G} are groups of variables.
- **The group Lasso with ℓ_∞ -norm:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_\infty$, where \mathcal{G} are groups of variables.
- **The sparse group Lasso:** same as above but with an additional ℓ_1 term.
- **The tree-structured sum of ℓ_2 -norms:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_2$, where \mathcal{G} is a tree-structured set of groups [15], and the η_g are positive weights.
- **The tree-structured sum of ℓ_∞ -norms:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_\infty$. See [15]
- **General sum of ℓ_∞ -norms:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g\|_\infty$, where no assumption are made on the groups \mathcal{G} .
- **The path-coding penalties** of [24].
- **the ℓ_1 -constraint.**
- **a few undocumented ones.**

Our software also handles regularization functions ψ on matrices \mathbf{W} in $\mathbb{R}^{p \times r}$ (note that \mathbf{W} can be transposed in these formulations). In particular,

- **The ℓ_1/ℓ_2 -norm:** $\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \|\mathbf{W}_i\|_2$, where \mathbf{W}_i denotes the i -th row of \mathbf{W} .
- **The ℓ_1/ℓ_∞ -norm:** $\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \|\mathbf{W}_i\|_\infty$,
- **The $\ell_1/\ell_2 + \ell_1$ -norm:** $\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \|\mathbf{W}_i\|_2 + \lambda_2 \sum_{i,j} |\mathbf{W}_{ij}|$.
- **The $\ell_1/\ell_\infty + \ell_1$ -norm:** $\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \|\mathbf{W}_i\|_\infty + \lambda_2 \sum_{i,j} |\mathbf{W}_{ij}|$,
- **The ℓ_1/ℓ_∞ -norm on rows and columns:** $\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \|\mathbf{W}_i\|_\infty + \lambda_2 \sum_{j=1}^r \|\mathbf{W}^j\|_\infty$, where \mathbf{W}^j denotes the j -th column of \mathbf{W} .
- **The multi-task tree-structured sum of ℓ_∞ -norms:**

$$\psi(\mathbf{W}) \triangleq \sum_{i=1}^p \sum_{g \in \mathcal{G}} \eta_g \|\mathbf{w}_g^i\|_\infty + \gamma \sum_{g \in \mathcal{G}} \eta_g \max_{j \in g} \|\mathbf{W}_j\|_\infty, \quad (25)$$

where the first double sums is in fact a sum of independent structured norms on the columns \mathbf{w}^i of \mathbf{W} , and the right term is a tree-structured regularization norm applied to the ℓ_∞ -norm of the rows of \mathbf{W} , thereby inducing the tree-structured regularization at the row level. \mathcal{G} is here a tree-structured set of groups.

- **The multi-task general sum of ℓ_∞ -norms** is the same as Eq. (25) except that the groups \mathcal{G} are general overlapping groups.
- **The trace norm:** $\psi(\mathbf{W}) \triangleq \|\mathbf{W}\|_*$.

Non-convex regularizations are also implemented with the ISTA algorithm (no duality gaps are of course provided in these cases):

- **The ℓ_0 -pseudo-norm:** $\psi(\mathbf{w}) \triangleq \|\mathbf{w}\|_0$.

- **The rank:** $\psi(\mathbf{W}) \triangleq \text{rank}(\mathbf{W})$.
- **The tree-structured ℓ_0 -pseudo-norm:** $\psi(\mathbf{w}) \triangleq \sum_{g \in \mathcal{G}} \delta_{\mathbf{w}_g \neq 0}$.

All of these regularization terms for vectors or matrices can be coupled with nonnegativity constraints. It is also possible to add an intercept, which one wishes not to regularize, and we will include this possibility in the next sections. There are also a few hidden undocumented options which are available in the source code.

We now present 3 functions for computing proximal operators associated to the previous regularization functions.

5.2 Function `spams.proximalFlat`

This function computes the proximal operators associated to many regularization functions, for input signals $\mathbf{U} = [\mathbf{u}^1, \dots, \mathbf{u}^n]$ in $\mathbb{R}^{p \times n}$, it finds a matrix $\mathbf{V} = [\mathbf{v}^1, \dots, \mathbf{v}^n]$ in $\mathbb{R}^{p \times n}$ such that:

- If one chooses a regularization function on vectors, for every column \mathbf{u} of \mathbf{U} , it computes one column \mathbf{v} of \mathbf{V} solving

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \|\mathbf{v}\|_0, \quad (26)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \|\mathbf{v}\|_1, \quad (27)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \frac{\lambda}{2} \|\mathbf{v}\|_2^2, \quad (28)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \|\mathbf{v}\|_1 + \lambda_2 \|\mathbf{v}\|_2^2, \quad (29)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{j=1}^{p-1} |\mathbf{v}_{j+1}^i - \mathbf{v}_j^i| + \lambda_2 \|\mathbf{v}\|_1 + \lambda_3 \|\mathbf{v}\|_2^2, \quad (30)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \delta^g(\mathbf{v}), \quad (31)$$

where \mathcal{T} is a tree-structured set of groups (see [16]), and $\delta^g(\mathbf{v}) = 0$ if $\mathbf{v}_g = 0$ and 1 otherwise. It can also solve

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \eta^g \|\mathbf{v}_g\|_2, \quad (32)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \eta^g \|\mathbf{v}_g\|_\infty, \quad (33)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{G}} \eta^g \|\mathbf{v}_g\|_\infty, \quad (34)$$

where \mathcal{G} is any kind of set of groups.

This function can also solve the following proximal operators on matrices

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{V}_i\|_2, \quad (35)$$

where \mathbf{V}_i is the i -th row of \mathbf{V} , or

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{V}_i\|_\infty, \quad (36)$$

or

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{V}_i\|_2 + \lambda_2 \sum_{i=1}^p \sum_{j=1}^n |\mathbf{V}_{ij}|, \quad (37)$$

or

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{V}_i\|_\infty + \lambda_2 \sum_{i=1}^p \sum_{j=1}^n |\mathbf{V}_{ij}|, \quad (38)$$

or

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{V}_i\|_\infty + \lambda_2 \sum_{j=1}^n \|\mathbf{V}^j\|_\infty. \quad (39)$$

where \mathbf{V}^j is the j -th column of \mathbf{V} .

See details below:

```
#
# Name: spams.proximalFlat
#
# Usage: spams.proximalFlat(U,return_val_loss = FALSE,numThreads =-1,lambda1=1.0,lambda2=0.,
#           lambda3=0.,intercept=FALSE,resetflow=FALSE,regul="",verbose=FALSE,
#           pos=FALSE,clever=TRUE,size_group=1,groups = NULL,transpose=FALSE)
#
# Description:
#   spams.proximalFlat computes proximal operators. Depending
#   on the value of regul, it computes
#
#   Given an input matrix U=[u^1,\ldots,u^n], it computes a matrix
#   V=[v^1,\ldots,v^n] such that
#   if one chooses a regularization functions on vectors, it computes
#   for each column u of U, a column v of V solving
#   if regul='l0'
#       argmin 0.5||u-v||_2^2 + lambda1||v||_0
#   if regul='l1'
#       argmin 0.5||u-v||_2^2 + lambda1||v||_1
#   if regul='l2'
#       argmin 0.5||u-v||_2^2 + 0.5lambda1||v||_2^2
#   if regul='elastic-net'
#       argmin 0.5||u-v||_2^2 + lambda1||v||_1 + lambda1_2||v||_2^2
#   if regul='fused-lasso'
#       argmin 0.5||u-v||_2^2 + lambda1 FL(v) + ...
#           ... lambda1_2||v||_1 + lambda1_3||v||_2^2
#   if regul='linf'
#       argmin 0.5||u-v||_2^2 + lambda1||v||_inf
#   if regul='l1-constraint'
#       argmin 0.5||u-v||_2^2 s.t. ||v||_1 <= lambda1
#   if regul='l2-not-squared'
#       argmin 0.5||u-v||_2^2 + lambda1||v||_2
#   if regul='group-lasso-l2'
#       argmin 0.5||u-v||_2^2 + lambda1 sum_g ||v_g||_2
#       where the groups are either defined by groups or by size_group,
#   if regul='group-lasso-linf'
#       argmin 0.5||u-v||_2^2 + lambda1 sum_g ||v_g||_inf
#   if regul='sparse-group-lasso-l2'
#       argmin 0.5||u-v||_2^2 + lambda1 sum_g ||v_g||_2 + lambda1_2 ||v||_1
#       where the groups are either defined by groups or by size_group,
#   if regul='sparse-group-lasso-linf'
#       argmin 0.5||u-v||_2^2 + lambda1 sum_g ||v_g||_inf + lambda1_2 ||v||_1
#   if regul='trace-norm-vec'
#       argmin 0.5||u-v||_2^2 + lambda1 ||mat(v)||_*
#       where mat(v) has size_group rows
```

```

#
#   if one chooses a regularization function on matrices
#   if regul='l1l2', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_{1/2}
#   if regul='l1linf', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_{1/inf}
#   if regul='l1l2+l1', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_{1/2} + lambda1_2||V||_{1/1}
#   if regul='l1linf+l1', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_{1/inf} + lambda1_2||V||_{1/1}
#   if regul='l1linf+row-column', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_{1/inf} + lambda1_2||V' ||_{1/inf}
#   if regul='trace-norm', V=
#       argmin 0.5||U-V||_F^2 + lambda1||V||_*
#   if regul='rank', V=
#       argmin 0.5||U-V||_F^2 + lambda1 rank(V)
#   if regul='none', V=
#       argmin 0.5||U-V||_F^2
#
#   for all these regularizations, it is possible to enforce non-negativity constraints
#   with the option pos, and to prevent the last row of U to be regularized, with
#   the option intercept
#
# Inputs:
#   U: double m x n matrix (input signals)
#       m is the signal size
#   return_val_loss:
#       if true the function will return 2 matrices in a list.
#   lambda1: (regularization parameter)
#   regul: (choice of regularization, see above)
#   lambda2: (optional, regularization parameter)
#   lambda3: (optional, regularization parameter)
#   verbose: (optional, verbosity level, false by default)
#   intercept: (optional, last row of U is not regularized,
#       false by default)
#   transpose: (optional, transpose the matrix in the regularization function)
#   size_group: (optional, for regularization functions assuming a group
#       structure). It is a scalar. When groups is not specified, it assumes
#       that the groups are the sets of consecutive elements of size size_group
#   groups: (int32, optional, for regularization functions assuming a group
#       structure. It is an int32 vector of size m containing the group indices of the
#       variables (first group is 1).
#   pos: (optional, adds positivity constraints on the
#       coefficients, false by default)
#   numThreads: (optional, number of threads for exploiting
#       multi-core / multi-cpus. By default, it takes the value -1,
#       which automatically selects all the available CPUs/cores).
#   resetflow: undocumented; modify at your own risks!
#   clever: undocumented; modify at your own risks!
#
# Output:
#   V: double m x n matrix (output coefficients)
#   val_regularizer: double 1 x n vector (value of the regularization
#   term at the optimum).

```

```

#     val_loss:         vector of size ncol(U)
#         alpha <- spams.proximalFlat(U,return_val_loss = FALSE, ...)
#         v <- spams.proximalFlat(U,return_val_loss = TRUE, ...)
#         alpha <- v[[1]]
#         val_loss <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#     Valid values for the regularization parameter (regul) are:
#     "l0", "l1", "l2", "l1nf", "l2-not-squared", "elastic-net", "fused-lasso",
#     "group-lasso-l2", "group-lasso-l1nf", "sparse-group-lasso-l2",
#     "sparse-group-lasso-l1nf", "l1l2", "l1l1nf", "l1l2+l1", "l1l1nf+l1",
#     "tree-l0", "tree-l2", "tree-l1nf", "graph", "graph-ridge", "graph-l2",
#     "multi-task-tree", "multi-task-graph", "l1l1nf-row-column", "trace-norm",
#     "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
m = 100;n = 1000
U = matrix(rnorm(m * n),nrow = m,ncol = n,byrow = FALSE)

# test l0
.printf("\nprox l0\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
  = 0.1,regul = 'l0', pos = FALSE,intercept = FALSE)),n = 1)

# test l1
.printf("\nprox l1, intercept, positivity constraint\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
  = 0.1,regul = 'l1', pos = TRUE,intercept = TRUE)),n = 1)

# test l2
.printf("\nprox squared-l2\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
  = 0.1,regul = 'l2', pos = FALSE,intercept = FALSE)),n = 1)

# test elastic-net
.printf("\nprox elastic-net\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
  = 0.1,regul = 'elastic-net', pos = FALSE,intercept = FALSE,lambda2 = 0.1)),n = 1)

# test fused-lasso
.printf("\nprox fused-lasso\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
  = 0.1,regul = 'fused-lasso', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n
  = 1)

# test l1l2
.printf("\nprox mixed norm l1/l2\n")

```

```

alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'l1l2', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n = 1)

# test l1linf
.printf("\nprox mixed norm l1/linf\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'l1linf', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n = 1)

# test l1l2+l1
.printf("\nprox mixed norm l1/l2 + l1\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'l1l2+l1', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n = 1)

# test l1linf+l1
.printf("\nprox mixed norm l1/linf + l1\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'l1linf+l1', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n =
    1)

# test l1linf-row-column
.printf("\nprox mixed norm l1/linf on rows and columns\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'l1linf-row-column', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 =
    0.1)),n = 1)

# test none
.printf("\nprox no regularization\n")
alpha = Xtest1('spams',quote(spams.proximalFlat(U,FALSE,numThreads = -1,verbose = TRUE,lambda1
    = 0.1,regul = 'none', pos = FALSE,intercept = FALSE,lambda2 = 0.1,lambda3 = 0.1)),n = 1)

```

5.3 Function spams.proximalTree

This function computes the proximal operators associated to tree-structured regularization functions, for input signals $\mathbf{U} = [\mathbf{u}^1, \dots, \mathbf{u}^n]$ in $\mathbb{R}^{p \times n}$, and a tree-structured set of groups [15], it computes a matrix $\mathbf{V} = [\mathbf{v}^1, \dots, \mathbf{v}^n]$ in $\mathbb{R}^{p \times n}$. When one uses a regularization function on vectors, it computes a column \mathbf{v} of \mathbf{V} for every column \mathbf{u} of \mathbf{U} :

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \eta^g \|\mathbf{v}_g\|_2, \quad (40)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \eta^g \|\mathbf{v}_g\|_\infty, \quad (41)$$

or

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{T}} \delta^g(\mathbf{v}), \quad (42)$$

where $\delta^g(\mathbf{v}) = 0$ if $\mathbf{v}_g = 0$ and 1 otherwise (see appendix of [16]).

When the multi-task tree-structured regularization function is used, it solves

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^n \sum_{g \in \mathcal{T}} \eta^g \|\mathbf{v}_g^i\|_\infty + \lambda_2 \sum_{g \in \mathcal{T}} \max_{j \in g} \|\mathbf{v}_g^j\|_\infty, \quad (43)$$

which is a formulation presented in [22].

This function can also be used for computing the proximal operators addressed by spams.proximalFlat (it will just not take into account the tree structure). The way the tree is incoded is presented below, (and examples are given in the file test_ProximalTree.m, with more usage details:

```

#
# Name: spams.proximalTree
#
# Usage: spams.proximalTree(U,tree,return_val_loss = FALSE,numThreads =-1,lambda1=1.0,lambda2
=0.,
#           lambda3=0.,intercept=FALSE,resetflow=FALSE,regul="",verbose=FALSE,
#           pos=FALSE,clever=TRUE,size_group=1,transpose=FALSE)
#
# Description:
#   spams.proximalTree computes a proximal operator. Depending
#   on the value of regul, it computes
#
#   Given an input matrix  $U=[u^1, \dots, u^n]$ , and a tree-structured set of groups  $T$ ,
#   it returns a matrix  $V=[v^1, \dots, v^n]$ :
#
#   when the regularization function is for vectors,
#   for every column  $u$  of  $U$ , it compute a column  $v$  of  $V$  solving
#   if regul='tree-l0'
#        $\operatorname{argmin} 0.5 \|u-v\|_2^2 + \lambda_{d1} \sum_{g \in T} \delta_g(v)$ 
#   if regul='tree-l2'
#       for all  $i$ ,  $v^i =$ 
#            $\operatorname{argmin} 0.5 \|u-v\|_2^2 + \lambda_{d1} \sum_{g \in T} \eta_g \|v_g\|_2$ 
#   if regul='tree-linf'
#       for all  $i$ ,  $v^i =$ 
#            $\operatorname{argmin} 0.5 \|u-v\|_2^2 + \lambda_{d1} \sum_{g \in T} \eta_g \|v_g\|_{\infty}$ 
#
#   when the regularization function is for matrices:
#   if regul='multi-task-tree'
#        $V = \operatorname{argmin} 0.5 \|U-V\|_F^2 + \lambda_{d1} \sum_{i=1}^n \sum_{g \in T} \eta_g \|v^i_g\|_{\infty} +$ 
#       ...
#            $\lambda_{d1_2} \sum_{g \in T} \eta_g \max_{\{j \in g\}} \|V_j\|_{\infty}$ 
#
#   it can also be used with any non-tree-structured regularization addressed by spams.
#   proximalFlat
#
#   for all these regularizations, it is possible to enforce non-negativity constraints
#   with the option pos, and to prevent the last row of  $U$  to be regularized, with
#   the option intercept
#
# Inputs:
#   U: double m x n matrix (input signals)
#       m is the signal size
#   tree: named list
#       with four fields, eta_g, groups, own_variables and N_own_variables.
#
#   The tree structure requires a particular organization of groups and variables
#   * Let us denote by  $N = |T|$ , the number of groups.
#       the groups should be ordered  $T=\{g_1, g_2, \dots, g_N\}$  such that if  $g_i$  is included
#       in  $g_j$ , then  $j \leq i$ .  $g_1$  should be the group at the root of the tree
#       and contains every variable.
#   * Every group is a set of contiguous indices for instance
#        $g_i=\{3,4,5\}$  or  $g_i=\{4,5,6,7\}$  or  $g_i=\{4\}$ , but not  $\{3,5\}$ ;
#   * We define  $\operatorname{root}(g_i)$  as the indices of the variables that are in  $g_i$ ,

```

```

#         but not in its descendants. For instance for
#         T={ g1={1,2,3,4},g2={2,3},g3={4} }, then, root(g1)={1},
#         root(g2)={2,3}, root(g3)={4},
#         We assume that for all i, root(gi) is a set of contiguous variables
#         * We assume that the smallest of root(gi) is also the smallest index of gi.
#
#         For instance,
#         T={ g1={1,2,3,4},g2={2,3},g3={4} }, is a valid set of groups.
#         but we can not have
#         T={ g1={1,2,3,4},g2={1,2},g3={3} }, since root(g1)={4} and 4 is not the
#         smallest element in g1.
#
#         We do not lose generality with these assumptions since they can be fulfilled for
any
#         tree-structured set of groups after a permutation of variables and a correct
ordering of the
#         groups.
#         see more examples in test_ProximalTree.m of valid tree-structured sets of groups.
#
#         The first fields sets the weights for every group
#         tree[['eta_g']]           double N vector
#
#         The next field sets inclusion relations between groups
#         (but not between groups and variables):
#         tree[['groups']]           sparse (double or boolean) N x N matrix
#         the (i,j) entry is non-zero if and only if i is different than j and
#         gi is included in gj.
#         the first column corresponds to the group at the root of the tree.
#
#         The next field define the smallest index of each group gi,
#         which is also the smallest index of root(gi)
#         tree[['own_variables']]    int32 N vector
#
#         The next field define for each group gi, the size of root(gi)
#         tree[['N_own_variables']]  int32 N vector
#
#         examples are given in test_ProximalTree.m
#
#         return_val_loss:
#         if true the function will return 2 matrices in a list.
#         lambda1: (regularization parameter)
#         regul: (choice of regularization, see above)
#         lambda2: (optional, regularization parameter)
#         lambda3: (optional, regularization parameter)
#         verbose: (optional, verbosity level, false by default)
#         intercept: (optional, last row of U is not regularized,
#         false by default)
#         pos: (optional, adds positivity constraints on the
#         coefficients, false by default)
#         transpose: (optional, transpose the matrix in the regularization function)
#         size_group: (optional, for regularization functions assuming a group
#         structure). It is a scalar. When groups is not specified, it assumes
#         that the groups are the sets of consecutive elements of size size_group
#         numThreads: (optional, number of threads for exploiting

```

```

# multi-core / multi-cpus. By default, it takes the value -1,
# which automatically selects all the available CPUs/cores).
# resetflow: undocumented; modify at your own risks!
# clever: undocumented; modify at your own risks!
#
# Output:
# V: double m x n matrix (output coefficients)
# val_regularizer: double 1 x n vector (value of the regularization
# term at the optimum).
# val_loss: vector of size ncol(U)
# alpha <- spams.proximalTree(U,tree,return_val_loss = FALSE, ...)
# v <- spams.proximalTree(U,tree,return_val_loss = TRUE, ...)
# alpha <- v[[1]]
# val_loss <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
# Valid values for the regularization parameter (regul) are:
# "l0", "l1", "l2", "linf", "l2-not-squared", "elastic-net", "fused-lasso",
# "group-lasso-l2", "group-lasso-linf", "sparse-group-lasso-l2",
# "sparse-group-lasso-linf", "l1l2", "l1linf", "l1l2+l1", "l1linf+l1",
# "tree-l0", "tree-l2", "tree-linf", "graph", "graph-ridge", "graph-l2",
# "multi-task-tree", "multi-task-graph", "l1linf-row-column", "trace-norm",
# "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
m = 10;n = 1000
U = matrix(rnorm(m * n),nrow = m,ncol = n,byrow = FALSE)

.printf("First tree example\n")
# Example 1 of tree structure
# tree structured groups:
# g1= {0 1 2 3 4 5 6 7 8 9}
# g2= {2 3 4}
# g3= {5 6 7 8 9}
own_variables = as.vector(c(0,2,5),mode= 'integer')
N_own_variables = as.vector(c(2,3,5),mode= 'integer')
eta_g = as.vector(c(1,1,1),mode = 'double')
groups = matrix(as.vector(c(0,0,0,
  1,0,0,
  1,0,0),mode='logical'),ncol = 3,byrow = T)
groups = as(groups, 'CsparseMatrix')
tree = list('eta_g'= eta_g, 'groups' = groups, 'own_variables' = own_variables,
  'N_own_variables' = N_own_variables)
.printf('\ntest prox tree-linf\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'tree-l2', pos = FALSE,intercept = FALSE)),n = 1)

```



```

.printf('\ntest prox tree-linf\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'tree-linf', pos = FALSE,intercept = FALSE)),n = 1)

.printf('Second tree example\n')
# Example 2 of tree structure
# tree structured groups:
# g1= {0 1 2 3 4 5 6 7 8 9}    root(g1) = { };
# g2= {0 1 2 3 4 5}          root(g2) = {0 1 2};
# g3= {3 4}                  root(g3) = {3 4};
# g4= {5}                    root(g4) = {5};
# g5= {6 7 8 9}              root(g5) = { };
# g6= {6 7}                  root(g6) = {6 7};
# g7= {8 9}                  root(g7) = {8};
# g8 = {9}                   root(g8) = {9};
own_variables = as.vector(c(0, 0, 3, 5, 6, 6, 8, 9),mode= 'integer')
N_own_variables = as.vector(c(0,3,2,1,0,2,1,1),mode= 'integer')
eta_g = as.vector(c(1,1,1,2,2,2,2.5,2.5),mode = 'double')
groups = matrix(as.vector(c(0,0,0,0,0,0,0,0,
  1,0,0,0,0,0,0,0,
  0,1,0,0,0,0,0,0,
  0,1,0,0,0,0,0,0,
  1,0,0,0,0,0,0,0,
  0,0,0,0,1,0,0,0,
  0,0,0,0,1,0,0,0,
  0,0,0,0,0,0,1,0),mode='logical'),ncol = 8,byrow = T)
groups = as(groups, 'CsparseMatrix')
tree = list('eta_g' = eta_g, 'groups' = groups, 'own_variables' = own_variables,
  'N_own_variables' = N_own_variables)
.printf('\ntest prox tree-l0\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'tree-l0', pos = FALSE,intercept = FALSE)),n = 1)

.printf('\ntest prox tree-l2\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'tree-l2', pos = FALSE,intercept = FALSE)),n = 1)

.printf('\ntest prox tree-linf\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'tree-linf', pos = FALSE,intercept = FALSE)),n = 1)

# mexProximalTree also works with non-tree-structured regularization functions
.printf('\nprox l1, intercept, positivity constraint\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,regul = 'l1', pos = TRUE,intercept = TRUE)),n = 1)

.printf('\nprox multi-task tree\n')
alpha = Xtest1('spams',quote(spams.proximalTree(U,tree,FALSE,numThreads = -1,verbose = TRUE,
  lambda1 = 0.1,lambda2 = 0.1,regul = 'multi-task-tree', pos = FALSE,intercept = FALSE)),n =
  1)

```

5.4 Function spams.proximalGraph

This function computes the proximal operators associated to structured sparse regularization, for input signals $\mathbf{U} = [\mathbf{u}^1, \dots, \mathbf{u}^n]$ in $\mathbb{R}^{p \times n}$, and a set of groups [22], it returns a matrix $\mathbf{V} = [\mathbf{v}^1, \dots, \mathbf{v}^n]$ in $\mathbb{R}^{p \times n}$. When one uses a regularization function on vectors, it computes a column \mathbf{v} of \mathbf{V} for every column \mathbf{u} of \mathbf{U} :

$$\min_{\mathbf{v} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 + \lambda \sum_{g \in \mathcal{G}} \eta^g \|\mathbf{v}_g\|_\infty, \quad (44)$$

or with a regularization function on matrices, it computes \mathbf{V} solving

$$\min_{\mathbf{V} \in \mathbb{R}^{p \times n}} \frac{1}{2} \|\mathbf{U} - \mathbf{V}\|_F^2 + \lambda \sum_{i=1}^n \sum_{g \in \mathcal{G}} \eta^g \|\mathbf{v}_g^i\|_\infty + \lambda_2 \sum_{g \in \mathcal{G}} \max_{j \in g} \|\mathbf{v}_g^j\|_\infty, \quad (45)$$

This function can also be used for computing the proximal operators addressed by spams.proximalFlat. The way the graph is incoded is presented below (and also in the example file test_ProximalGraph.m, with more usage details:

```
#
# Name: spams.proximalGraph
#
# Usage: spams.proximalGraph(U,graph,return_val_loss = FALSE,numThreads = -1,lambda1=1.0,lambda2
# =0.,
#         lambda3=0.,intercept=FALSE,resetflow=FALSE,regul="",verbose=FALSE,
#         pos=FALSE,clever=TRUE,size_group=1,transpose=FALSE)
#
# Description:
#   spams.proximalGraph computes a proximal operator. Depending
#   on the value of regul, it computes
#
#   Given an input matrix U=[u^1,\ldots,u^n], and a set of groups G,
#   it computes a matrix V=[v^1,\ldots,v^n] such that
#
#   if regul='graph'
#   for every column u of U, it computes a column v of V solving
#       argmin 0.5||u-v||_2^2 + lambda1\sum_{g \in G} \eta_g||v_g||_inf
#
#   if regul='graph+ridge'
#   for every column u of U, it computes a column v of V solving
#       argmin 0.5||u-v||_2^2 + lambda1\sum_{g \in G} \eta_g||v_g||_inf + lambda1_2||v||_
2^2
#
#   if regul='multi-task-graph'
#       V=argmin 0.5||U-V||_F^2 + lambda1 \sum_{i=1}^n\sum_{g \in G} \eta_g||v^i_g||_inf +
...
#                               lambda1_2 \sum_{g \in G} \eta_g max_{j in g}||
V_j||_inf}
#
#   it can also be used with any regularization addressed by spams.proximalFlat
#
#   for all these regularizations, it is possible to enforce non-negativity constraints
#   with the option pos, and to prevent the last row of U to be regularized, with
#   the option intercept
#
# Inputs:
```

```

# U: double p x n matrix (input signals)
#     m is the signal size
# graph: struct
#     with three fields, eta_g, groups, and groups_var
#
#     The first fields sets the weights for every group
#     graph.eta_g         double N vector
#
#     The next field sets inclusion relations between groups
#     (but not between groups and variables):
#     graph.groups        sparse (double or boolean) N x N matrix
#     the (i,j) entry is non-zero if and only if i is different than j and
#     gi is included in gj.
#
#     The next field sets inclusion relations between groups and variables
#     graph.groups_var    sparse (double or boolean) p x N matrix
#     the (i,j) entry is non-zero if and only if the variable i is included
#     in gj, but not in any children of gj.
#
#     examples are given in test_ProximalGraph.m
#
# return_val_loss:
#     if true the function will return 2 matrices in a list.
# lambda1: (regularization parameter)
# regul: (choice of regularization, see above)
# lambda2: (optional, regularization parameter)
# lambda3: (optional, regularization parameter)
# verbose: (optional, verbosity level, false by default)
# intercept: (optional, last row of U is not regularized,
#     false by default)
# pos: (optional, adds positivity constraints on the
#     coefficients, false by default)
# numThreads: (optional, number of threads for exploiting
#     multi-core / multi-cpus. By default, it takes the value -1,
#     which automatically selects all the available CPUs/cores).
# resetflow:    undocumented; modify at your own risks!
# clever:        undocumented; modify at your own risks!
# size_group:   undocumented; modify at your own risks!
# transpose:    undocumented; modify at your own risks!
#
# Output:
# V: double p x n matrix (output coefficients)
# val_regularizer: double 1 x n vector (value of the regularization
# term at the optimum).
# val_loss:      vector of size ncol(U)
#     alpha <- spams.proximalGraph(U,graph,return_val_loss = FALSE, ...)
#     v <- spams.proximalGraph(U,graph,return_val_loss = TRUE, ...)
#     alpha <- v[[1]]
#     val_loss <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

```

# Note:
#   Valid values for the regularization parameter (regul) are:
#   "l0", "l1", "l2", "l1f", "l2-not-squared", "elastic-net", "fused-lasso",
#   "group-lasso-l2", "group-lasso-l1f", "sparse-group-lasso-l2",
#   "sparse-group-lasso-l1f", "l1l2", "l1l1f", "l1l2+l1", "l1l1f+l1",
#   "tree-l0", "tree-l2", "tree-l1f", "graph", "graph-ridge", "graph-l2",
#   "multi-task-tree", "multi-task-graph", "l1l1f-row-column", "trace-norm",
#   "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
lambda1 = 0.1 # regularization parameter
num_threads = -1 # all cores (-1 by default)
verbose = TRUE # verbosity, false by default
pos = FALSE # can be used with all the other regularizations
intercept = FALSE # can be used with all the other regularizations

U = matrix(rnorm(1000),nrow = 10,ncol = 100,byrow = FALSE)
.printf('First graph example\n')
# Example 1 of graph structure
# groups:
# g1= {0 1 2 3}
# g2= {3 4 5 6}
# g3= {6 7 8 9}
eta_g = as.vector(c(1, 1, 1),mode='double')
groups = as(matrix(as.vector(rep(0,9),mode='logical'),nrow = 3,ncol = 3),'CsparseMatrix')
groups_var = as(matrix(as.vector(c(1, 0, 0,
  1, 0, 0,
  1, 0, 0,
  1, 1, 0,
  0, 1, 0,
  0, 1, 0,
  0, 1, 1,
  0, 0, 1,
  0, 0, 1,
  0, 0, 1),mode='logical'),ncol = 3,byrow = T),'CsparseMatrix')
graph = list('eta_g' = eta_g, 'groups' = groups, 'groups_var' = groups_var)
.printf('\ntest prox graph\n')
regul='graph'
alpha = Xtest1('spams',quote(spams.proximalGraph(U,graph,FALSE,lambda1 = lambda1,numThreads =
  num_threads ,verbose = verbose,pos = pos,intercept = intercept,regul = regul)),n = 1)

# Example 2 of graph structure
# groups:
# g1= {0 1 2 3}
# g2= {3 4 5 6}
# g3= {6 7 8 9}
# g4= {0 1 2 3 4 5}
# g5= {6 7 8}

eta_g = as.vector(c(1, 1, 1, 1, 1),mode='double')

```

```

groups = as(matrix(as.vector(c(0, 0, 0, 1, 0,
0, 0, 0, 0, 0,
0, 0, 0, 0, 0,
0, 0, 0, 0, 0,
0, 0, 1, 0, 0)),mode='logical'),ncol = 5,byrow = T),'CsparseMatrix')

groups_var = as(matrix(as.vector(c(1, 0, 0, 0, 0,
1, 0, 0, 0, 0,
1, 0, 0, 0, 0,
1, 1, 0, 0, 0,
0, 1, 0, 1, 0,
0, 1, 0, 1, 0,
0, 1, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 1, 0, 0)),mode='logical'),ncol = 5,byrow = T),'CsparseMatrix')
graph = list('eta_g'= eta_g, 'groups' = groups, 'groups_var' = groups_var)

.printf('\ntest prox graph\n')
alpha = Xtest1('spams',quote(spams.proximalGraph(U,graph,FALSE,lambda1 = lambda1,numThreads =
num_threads ,verbose = verbose,pos = pos,intercept = intercept,regul = regul)),n = 1)
#
.printf('\ntest prox multi-task-graph\n')
regul = 'multi-task-graph'
lambda2 = 0.1
alpha = Xtest1('spams',quote(spams.proximalGraph(U,graph,FALSE,lambda1 = lambda1,lambda2 =
lambda2,numThreads = num_threads ,verbose = verbose,pos = pos,intercept = intercept,regul
= regul)),n = 1)
#
.printf('\ntest no regularization\n')
regul = 'none'
alpha = Xtest1('spams',quote(spams.proximalGraph(U,graph,FALSE,lambda1 = lambda1,lambda2 =
lambda2,numThreads = num_threads ,verbose = verbose,pos = pos,intercept = intercept,regul
= regul)),n = 1)

```

5.5 Function spams.proximalPathCoding

This function computes the proximal operators associated to the path coding penalties of [24].

```

#
# The R function is not yet implemented.
#

```

The following piece of code illustrates how to use this function. This function is associated to a function to evaluate the penalties:

5.6 Function spams.evalPathCoding

```

#
# The R function is not yet implemented.
#

```

The following piece of code illustrates how to use this function.

After having presented the regularization terms which our software can handle, we present the various formulations that we address

5.7 Problems Addressed

We present here regression or classification formulations and their multi-task variants.

5.7.1 Regression Problems with the Square Loss

Given a training set $\{\mathbf{x}^i, y_i\}_{i=1}^n$, with $\mathbf{x}^i \in \mathbb{R}^p$ and $y_i \in \mathbb{R}$ for all i in $[1; n]$, we address

$$\min_{\mathbf{w} \in \mathbb{R}^p, b \in \mathbb{R}} \sum_{i=1}^n \frac{1}{2} (y_i - \mathbf{w}^\top \mathbf{x}^i - b)^2 + \lambda \psi(\mathbf{w}),$$

where b is an optional variable acting as an ‘‘intercept’’, which is not regularized, and ψ can be any of the regularization functions presented above. Let us consider the vector \mathbf{y} in \mathbb{R}^n that carries the entries y_i . The problem without the intercept takes the following form, which we have already encountered in the previous toolbox, but with different notations:

$$\min_{\mathbf{w} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \psi(\mathbf{w}),$$

where the $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]^\top$ (the \mathbf{x}^i 's are here the rows of \mathbf{X}).

5.7.2 Classification Problems with the Logistic Loss

The next formulation that our software can solve is the regularized logistic regression formulation. We are again given a training set $\{\mathbf{x}^i, y_i\}_{i=1}^n$, with $\mathbf{x}^i \in \mathbb{R}^p$, but the variables y_i are now in $\{-1, +1\}$ for all i in $[1; n]$. The optimization problem we address is

$$\min_{\mathbf{w} \in \mathbb{R}^p, b \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \log(1 + e^{-y_i(\mathbf{w}^\top \mathbf{x}^i + b)}) + \lambda \psi(\mathbf{w}),$$

with again ψ taken to be one of the regularization function presented above, and b is an optional intercept.

5.7.3 Multi-class Classification Problems with the Softmax Loss

We have also implemented a multi-class logistic classifier (or softmax). For a classification problem with r classes, we are given a training set $\{\mathbf{x}^i, y_i\}_{i=1}^n$, where the variables \mathbf{x}^i are still vectors in \mathbb{R}^p , but the y_i 's have integer values in $\{1, 2, \dots, r\}$. The formulation we address is the following multi-class learning problem

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}, \mathbf{b} \in \mathbb{R}^r} \frac{1}{n} \sum_{i=1}^n \log \left(\sum_{j=1}^r e^{(\mathbf{w}^j - \mathbf{w}^{y_i})^\top \mathbf{x}^i + \mathbf{b}_j - \mathbf{b}_{y_i}} \right) + \lambda \sum_{j=1}^r \psi(\mathbf{w}^j), \quad (46)$$

where $\mathbf{W} = [\mathbf{w}^1, \dots, \mathbf{w}^r]$ and the optional vector \mathbf{b} in \mathbb{R}^r carries intercepts for each class.

5.7.4 Multi-task Regression Problems with the Square Loss

We are now considering a problem with r tasks, and a training set $\{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^n$, where the variables \mathbf{x}^i are still vectors in \mathbb{R}^p , and \mathbf{y}^i is a vector in \mathbb{R}^r . We are looking for r regression vectors \mathbf{w}^j , for $j \in [1; r]$, or equivalently for a matrix $\mathbf{W} = [\mathbf{w}^1, \dots, \mathbf{w}^r]$ in $\mathbb{R}^{p \times r}$. The formulation we address is the following multi-task regression problem

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}, \mathbf{b} \in \mathbb{R}^r} \sum_{j=1}^r \sum_{i=1}^n \frac{1}{2} (\mathbf{y}_j^i - \mathbf{w}^j \mathbf{x}^i - \mathbf{b}_j)^2 + \lambda \psi(\mathbf{W}),$$

where ψ is any of the regularization function on matrices we have presented in the previous section. Note that by introducing the appropriate variables \mathbf{Y} , the problem without intercept could be equivalently rewritten

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}} \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\mathbf{W}\|_F^2 + \lambda \psi(\mathbf{W}).$$

5.7.5 Multi-task Classification Problems with the Logistic Loss

The multi-task version of the logistic regression follows the same principle. We consider r tasks, and a training set $\{\mathbf{x}^i, \mathbf{y}^i\}_{i=1}^n$, with the \mathbf{x}^i 's in \mathbb{R}^p , and the \mathbf{y}^i 's are vectors in $\{-1, +1\}^r$. We look for a matrix $\mathbf{W} = [\mathbf{w}^1, \dots, \mathbf{w}^r]$ in $\mathbb{R}^{p \times r}$. The formulation is the following multi-task regression problem

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}, \mathbf{b} \in \mathbb{R}^r} \sum_{j=1}^r \frac{1}{n} \sum_{i=1}^n \log \left(1 + e^{-\mathbf{y}_j^i (\mathbf{w}^j \top \mathbf{x}^i + \mathbf{b}_j)} \right) + \lambda \psi(\mathbf{W}).$$

5.7.6 Multi-task and Multi-class Classification Problems with the Softmax Loss

The multi-task/multi-class version directly follows from the formulation of Eq. (46), but associates with each class a task, and as a consequence, regularizes the matrix \mathbf{W} in a particular way:

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times r}, \mathbf{b} \in \mathbb{R}^r} \frac{1}{n} \sum_{i=1}^n \log \left(\sum_{j=1}^r e^{(\mathbf{w}^j - \mathbf{w}^{\mathbf{y}_i}) \top \mathbf{x}^i + \mathbf{b}_j - \mathbf{b}_{\mathbf{y}_i}} \right) + \lambda \psi(\mathbf{W}).$$

How duality gaps are computed for any of these formulations is presented in Appendix A. We now present the main functions for solving these problems

5.8 Function spams.fistaFlat

Given a matrix $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^p]^T$ in $\mathbb{R}^{m \times p}$, and a matrix $\mathbf{Y} = [\mathbf{y}^1, \dots, \mathbf{y}^n]$, it solves the optimization problems presented in the previous section, with the same regularization functions as spams.proximalFlat. see usage details below:

```
#
# Name: spams.fistaFlat
#
# Usage: spams.fistaFlat(Y,X,W0,return_optim_info = FALSE,numThreads =-1,max_it =1000,L0=1.0,
#         fixed_step=FALSE,gamma=1.5,lambda1=1.0,delta=1.0,lambda2=0.,lambda3=0.,
#         a=1.0,b=0.,c=1.0,tol=0.000001,it0=100,max_iter_backtracking=1000,
#         compute_gram=FALSE,lin_admm=FALSE,admm=FALSE,intercept=FALSE,
#         resetflow=FALSE,regul="",loss="",verbose=FALSE,pos=FALSE,clever=FALSE,
#         log=FALSE,ista=FALSE,subgrad=FALSE,logName="",is_inner_weights=FALSE,
#         inner_weights=c(.),size_group=1,groups = NULL,sqrt_step=TRUE,
#         transpose=FALSE,linesearch_mode=0)
#
# Description:
#   spams.fistaFlat solves sparse regularized problems.
#   X is a design matrix of size m x p
#   X=[x^1,...,x^n]', where the x_i's are the rows of X
#   Y=[y^1,...,y^n] is a matrix of size m x n
#   It implements the algorithms FISTA, ISTA and subgradient descent.
#
#   - if loss='square' and regul is a regularization function for vectors,
#     the entries of Y are real-valued, W = [w^1,...,w^n] is a matrix of size p x n
#     For all column y of Y, it computes a column w of W such that
#       w = argmin 0.5||y- X w||_2^2 + lambda1 psi(w)
#
#   - if loss='square' and regul is a regularization function for matrices
#     the entries of Y are real-valued, W is a matrix of size p x n.
#     It computes the matrix W such that
#       W = argmin 0.5||Y- X W||_F^2 + lambda1 psi(W)
#
#   - loss='square-missing' same as loss='square', but handles missing data
#     represented by NaN (not a number) in the matrix Y
```

```

#
#   - if loss='logistic' and regul is a regularization function for vectors,
#   the entries of Y are either -1 or +1,  $W = [w^1, \dots, w^n]$  is a matrix of size  $p \times n$ 
#   For all column  $y$  of Y, it computes a column  $w$  of W such that
#        $w = \operatorname{argmin} (1/m) \sum_{j=1}^m \log(1 + e^{(-y_j x^j' w)}) + \lambda_1 \operatorname{psi}(w)$ ,
#   where  $x^j$  is the  $j$ -th row of X.
#
#   - if loss='logistic' and regul is a regularization function for matrices
#   the entries of Y are either -1 or +1, W is a matrix of size  $p \times n$ 
#        $W = \operatorname{argmin} \sum_{i=1}^n (1/m) \sum_{j=1}^m \log(1 + e^{(-y^i_j x^j' w^i)}) + \lambda_1 \operatorname{psi}$ 
(W)
#
#   - if loss='multi-logistic' and regul is a regularization function for vectors,
#   the entries of Y are in  $\{0, 1, \dots, N\}$  where N is the total number of classes
#    $W = [W^1, \dots, W^n]$  is a matrix of size  $p \times Nn$ , each submatrix  $W^i$  is of size  $p \times N$ 
#   for all submatrix  $WW$  of W, and column  $y$  of Y, it computes
#        $WW = \operatorname{argmin} (1/m) \sum_{j=1}^m \log(\sum_{j=1}^r e^{(x^j' (ww^j - ww^{\{y_j\}}))}) + \lambda_1$ 
sum_{j=1}^N \operatorname{psi}(ww^j),
#   where  $ww^j$  is the  $j$ -th column of WW.
#
#   - if loss='multi-logistic' and regul is a regularization function for matrices,
#   the entries of Y are in  $\{0, 1, \dots, N\}$  where N is the total number of classes
#   W is a matrix of size  $p \times N$ , it computes
#        $W = \operatorname{argmin} (1/m) \sum_{j=1}^m \log(\sum_{j=1}^r e^{(x^j' (w^j - w^{\{y_j\}}))}) + \lambda_1$ 
psi(W)
#   where  $ww^j$  is the  $j$ -th column of WW.
#
#   - loss='cur' useful to perform sparse CUR matrix decompositions,
#        $W = \operatorname{argmin} 0.5 ||Y - X * W * X||_F^2 + \lambda_1 \operatorname{psi}(W)$ 
#
#
#   The function psi are those used by spams.proximalFlat (see documentation)
#
#   This function can also handle intercepts (last row of W is not regularized),
#   and/or non-negativity constraints on W, and sparse matrices for X
#
# Inputs:
#   Y: double dense  $m \times n$  matrix
#   X: double dense or sparse  $m \times p$  matrix
#   W0: double dense  $p \times n$  matrix or  $p \times Nn$  matrix (for multi-logistic loss)
#       initial guess
#   return_optim_info:
#       if true the function will return 2 matrices in a list.
#   loss: (choice of loss, see above)
#   regul: (choice of regularization, see function spams.proximalFlat)
#   lambda1: (regularization parameter)
#   lambda2: (optional, regularization parameter, 0 by default)
#   lambda3: (optional, regularization parameter, 0 by default)
#   verbose: (optional, verbosity level, false by default)
#   pos: (optional, adds positivity constraints on the
#       coefficients, false by default)
#   transpose: (optional, transpose the matrix in the regularization function)
#   size_group: (optional, for regularization functions assuming a group
#       structure)

```



```

# groups: (int32, optional, for regularization functions assuming a group
#         structure, see spams.proximalFlat)
# numThreads: (optional, number of threads for exploiting
#             multi-core / multi-cpus. By default, it takes the value -1,
#             which automatically selects all the available CPUs/cores).
# max_it: (optional, maximum number of iterations, 100 by default)
# it0: (optional, frequency for computing duality gap, every 10 iterations by default)
# tol: (optional, tolerance for stopping criterion, which is a relative duality gap
#      if it is available, or a relative change of parameters).
# gamma: (optional, multiplier for increasing the parameter L in fista, 1.5 by default)
# L0: (optional, initial parameter L in fista, 0.1 by default, should be small enough)
# fixed_step: (deactivate the line search for L in fista and use L0 instead)
# linesearch_mode: (line-search scheme when ista=true:
# 0: default, monotonic backtracking scheme
# 1: monotonic backtracking scheme, with restart at each iteration
# 2: Barzilai-Borwein step sizes (similar to SparSA by Wright et al.)
# 3: non-monotonic backtracking
# compute_gram: (optional, pre-compute  $X^T X$ , false by default).
# intercept: (optional, do not regularize last row of W, false by default).
# ista: (optional, use ista instead of fista, false by default).
# subgrad: (optional, if not ista, use subgradient descent instead of fista, false by
default).
# a:
# b: (optional, if subgrad, the gradient step is  $a/(t+b)$ 
# also similar options as spams.proximalFlat
#
# the function also implements the ADMM algorithm via an option admm=true. It is not
documented
# and you need to look at the source code to use it.
# delta:  undocumented; modify at your own risks!
# c:      undocumented; modify at your own risks!
# max_iter_backtracking:  undocumented; modify at your own risks!
# lin_admm:  undocumented; modify at your own risks!
# admm:      undocumented; modify at your own risks!
# resetflow:  undocumented; modify at your own risks!
# clever:     undocumented; modify at your own risks!
# log:       undocumented; modify at your own risks!
# logName:   undocumented; modify at your own risks!
# is_inner_weights:  undocumented; modify at your own risks!
# inner_weights:  undocumented; modify at your own risks!
# sqrt_step:   undocumented; modify at your own risks!
#
# Output:
# W: double dense p x n matrix or p x Nn matrix (for multi-logistic loss)
# optim: optional, double dense 4 x n matrix.
#       first row: values of the objective functions.
#       third row: values of the relative duality gap (if available)
#       fourth row: number of iterations
# optim_info: vector of size 4, containing information of the optimization.
#       W <- spams.fistaFlat(Y,X,W0,return_optim_info = FALSE,...)
#       v <- spams.fistaFlat(Y,X,W0,return_optim_info = TRUE,...)
#       W <- v[[1]]
#       optim_info <- v[[2]]
#

```

```

# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   Valid values for the regularization parameter (regul) are:
#   "l0", "l1", "l2", "l1nf", "l2-not-squared", "elastic-net", "fused-lasso",
#   "group-lasso-l2", "group-lasso-l1nf", "sparse-group-lasso-l2",
#   "sparse-group-lasso-l1nf", "l1l2", "l1l1nf", "l1l2+l1", "l1l1nf+l1",
#   "tree-l0", "tree-l2", "tree-l1nf", "graph", "graph-ridge", "graph-l2",
#   "multi-task-tree", "multi-task-graph", "l1l1nf-row-column", "trace-norm",
#   "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
m = 100;n = 200
X = matrix(rnorm(m * n),nrow = m,ncol = n,byrow = FALSE)
X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = spams.normalize(X)
Y = matrix(rnorm(m),nrow = m,ncol = 1,byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
# Regression experiments
# 100 regression problems with the same design matrix X.
.printf("\nVarious regression experiments\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,compute_
  gram = TRUE, loss = 'square',regul = 'l1')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])
###

.printf("\nFISTA + Regression l1\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,compute_
  gram = TRUE, loss = 'square',regul = 'l1',ista = TRUE)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])
###

.printf("\nSubgradient Descent + Regression l1\n")
max_it = 200
it0 = 10
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = 50, max_it = 500,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,compute_
  gram = TRUE, loss = 'square',regul = 'l1',ista = FALSE,subgrad = TRUE,a = 0.1, b = 100)),n
  = 1)

```

```

W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
      [1],optim_info[3],optim_info[4])
###
.printf("\nFISTA + Regression l2\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
      0.05, it0 = 50, max_it = 500,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,compute_
      gram = TRUE, loss = 'square',regul = 'l2',ista = FALSE,subgrad = TRUE,a = 0.1, b = 1000)),n
      = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
      [1],optim_info[3],optim_info[4])
###
.printf("\nFISTA + Regression l2 + sparse feature matrix\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,as(X,'CsparseMatrix'),W0,TRUE,numThreads = 1,
      verbose = TRUE,lambda1 = 0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept =
      FALSE,pos = FALSE,compute_gram = TRUE, loss = 'square',regul = 'l2',ista = FALSE,subgrad =
      FALSE,a = 0.1, b = 1000)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
      [1],optim_info[3],optim_info[4])
#####
.printf("\nFISTA + Regression Elastic-Net\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
      0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
      compute_gram = TRUE, loss = 'square',regul = 'elastic-net',ista = FALSE,subgrad = FALSE,a =
      0.1, b = 1000,lambda2 = 0.1)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
      [1],optim_info[3],optim_info[4])

.printf("\nFISTA + Group Lasso L2\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
      0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
      compute_gram = TRUE, loss = 'square',regul = 'group-lasso-l2',ista = FALSE,subgrad = FALSE,
      a = 0.1, b = 1000,lambda2 = 0.1,size_group = 2)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nFISTA + Group Lasso L2 with variable size of groups\n")
groups = matrix(sample(1:5,ncol(X),replace=TRUE),nrow = 1)
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
      0.5, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,compute
      _gram = TRUE, loss = 'square',regul = 'group-lasso-l2',ista = FALSE,subgrad = FALSE,a =
      0.1, b = 1000,lambda2 = 0.1,size_group = 2,
      groups = groups)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

```

```

.printf("\nFISTA + Trace Norm\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = TRUE, loss = 'square',regul = 'trace-norm-vec',ista = FALSE,subgrad = FALSE,
  a = 0.1, b = 1000,lambda2 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

###
.printf("\nFISTA + Regression Fused-Lasso\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = TRUE, loss = 'square',regul = 'fused-lasso',ista = FALSE,subgrad = FALSE,a =
  0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nFISTA + Regression no regularization\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = TRUE, loss = 'square',regul = 'none',ista = FALSE,subgrad = FALSE,a = 0.1, b
  = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nFISTA + Regression l1 with intercept\n")
x1 = cbind(X,matrix(1,nrow = nrow(X),ncol = 1))
W01 = rbind(W0,matrix(0,nrow = 1, ncol = ncol(W0)))
res = Xtest1('spams',quote(spams.fistaFlat(Y,x1,W01,TRUE,numThreads = 1,verbose = TRUE,lambda1
  = 0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = TRUE,pos = FALSE,
  compute_gram = TRUE, loss = 'square',regul = 'l1',ista = FALSE,subgrad = FALSE,a = 0.1, b =
  1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

.printf("\nFISTA + Regression l1 with intercept+ non-negative\n")
x1 = cbind(X,matrix(1,nrow = nrow(X),ncol = 1))
W01 = rbind(W0,matrix(0,nrow = 1, ncol = ncol(W0)))
res = Xtest1('spams',quote(spams.fistaFlat(Y,x1,W01,TRUE,numThreads = 1,verbose = TRUE,lambda1
  = 0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = TRUE,pos = TRUE,
  compute_gram = TRUE, loss = 'square',regul = 'l1',ista = FALSE,subgrad = FALSE,a = 0.1, b =
  1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nISTA + Regression l0\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.05, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,

```

```

compute_gram = TRUE, loss = 'square', regul = 'l0', ista = FALSE, subgrad = FALSE, a = 0.1, b =
  1000, lambda2 = 0.1, lambda3 = 0.1, size_group = 5)), n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n", optim_info[1], optim_info[4])
# Classification

.printf("\nOne classification experiment\n")
Y = matrix(2. * as.double(rnorm(100) > 0.) - 1., nrow = 100, ncol = 1, byrow = FALSE)
.printf("\nFISTA + Logistic l1\n")
res = Xtest1('spams', quote(spams.fistaFlat(Y, X, W0, TRUE, numThreads = 1, verbose = TRUE, lambda1 =
  0.01, it0 = it0, max_it = max_it, L0 = 0.1, tol = 1e-3, intercept = FALSE, pos = FALSE,
  compute_gram = TRUE, loss = 'logistic', regul = 'l1', ista = FALSE, subgrad = FALSE, a = 0.1, b
  = 1000, lambda2 = 0.1, lambda3 = 0.1, size_group = 5)), n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n", optim_info
  [1], optim_info[3], optim_info[4])
# can be used of course with other regularization functions, intercept,...
res = Xtest1('spams', quote(spams.fistaFlat(Y, X, W0, TRUE, numThreads = 1, verbose = TRUE, lambda1 =
  0.01, it0 = it0, max_it = max_it, L0 = 0.1, tol = 1e-3, intercept = FALSE, pos = FALSE,
  compute_gram = TRUE, loss = 'weighted-logistic', regul = 'l1', ista = FALSE, subgrad = FALSE, a
  = 0.1, b = 1000, lambda2 = 0.1, lambda3 = 0.1, size_group = 5)), n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n", optim_info
  [1], optim_info[3], optim_info[4])
# can be used of course with other regularization functions, intercept,...

.printf("\nFISTA + Logistic l1 + sparse matrix\n")
res = Xtest1('spams', quote(spams.fistaFlat(Y, as(X, 'CsparseMatrix'), W0, TRUE, numThreads = 1,
  verbose = TRUE, lambda1 = 0.01, it0 = it0, max_it = max_it, L0 = 0.1, tol = 1e-3, intercept =
  FALSE, pos = FALSE, compute_gram = TRUE, loss = 'logistic', regul = 'l1', ista = FALSE, subgrad
  = FALSE, a = 0.1, b = 1000, lambda2 = 0.1, lambda3 = 0.1, size_group = 5)), n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n", optim_info
  [1], optim_info[3], optim_info[4])
# can be used of course with other regularization functions, intercept,...

# Multi-Class classification
Y = ceiling(5 * matrix(runif(100 * 1000, 0, 1), nrow = 100, ncol = 1000, byrow = FALSE)) - 1
.printf("\nFISTA + Multi-Class Logistic l1\n")
nclasses = max(Y) + 1
W0 = matrix(0, nrow = ncol(X), nclasses * ncol(Y), byrow = FALSE)
res = Xtest1('spams', quote(spams.fistaFlat(Y, X, W0, TRUE, numThreads = 1, verbose = TRUE, lambda1 =
  0.01, it0 = it0, max_it = max_it, L0 = 0.1, tol = 1e-3, intercept = FALSE, pos = FALSE,
  compute_gram = TRUE, loss = 'multi-logistic', regul = 'l1', ista = FALSE, subgrad = FALSE, a =
  0.1, b = 1000, lambda2 = 0.1, lambda3 = 0.1, size_group = 5)), n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n", optim_info
  [1], optim_info[3], optim_info[4])
# can be used of course with other regularization functions, intercept,...

```

```

# Multi-Task regression
Y = matrix(rnorm(100 * 100),nrow = 100,ncol = 100,byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
.printf("\nFISTA + Regression l1l2\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'square',regul = 'l1l2',ista = FALSE,subgrad = FALSE,a = 0.1,
  b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

.printf("\nFISTA + Regression l1linf\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'square',regul = 'l1linf',ista = FALSE,subgrad = FALSE,a =
  0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

.printf("\nFISTA + Regression l1l2 + l1\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'square',regul = 'l1l2+l1',ista = FALSE,subgrad = FALSE,a =
  0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nFISTA + Regression l1linf + l1\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'square',regul = 'l1linf+l1',ista = FALSE,subgrad = FALSE,a =
  0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

.printf("\nFISTA + Regression l1linf + row + columns\n")
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'square',regul = 'l1linf-row-column',ista = FALSE,subgrad =
  FALSE,a = 0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

```



```

# Multi-Task Classification

.printf("\nFISTA + Logistic + l1l2\n")
Y = matrix(2. * as.double(rnorm(100 * 100) > 1.) - 1.,nrow = 100,ncol = 100,byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'logistic',regul = 'l1l2',ista = FALSE,subgrad = FALSE,a =
  0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])
# Multi-Class + Multi-Task Regularization

.printf("\nFISTA + Multi-Class Logistic l1l2\n")
Y = ceiling(5 * matrix(runif(100 * 1000,0,1),nrow = 100,ncol = 1000,byrow = FALSE)) - 1
Y = spams.normalize(Y)
nclasses = max(Y) + 1
W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaFlat(Y,X,W0,TRUE,numThreads = 1,verbose = TRUE,lambda1 =
  0.01, it0 = it0, max_it = max_it,L0 = 0.1, tol = 1e-3, intercept = FALSE,pos = FALSE,
  compute_gram = FALSE, loss = 'multi-logistic',regul = 'l1l2',ista = FALSE,subgrad = FALSE,a
  = 0.1, b = 1000,lambda2 = 0.1,lambda3 = 0.1,size_group = 5)),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])
# can be used of course with other regularization functions, intercept,...

```

5.9 Function spams.fistaTree

Given a matrix $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^p]^T$ in $\mathbb{R}^{m \times p}$, and a matrix $\mathbf{Y} = [\mathbf{y}^1, \dots, \mathbf{y}^n]$, it solves the optimization problems presented in the previous section, with the same regularization functions as `spams.proximalTree`. see usage details below:

```

#
# Name: spams.fistaTree
#
# Usage: spams.fistaTree(Y,X,W0,tree,return_optim_info = FALSE,numThreads =-1,max_it =1000,L0
  =1.0,
#         fixed_step=FALSE,gamma=1.5,lambda1=1.0,delta=1.0,lambda2=0.,lambda3=0.,
#         a=1.0,b=0.,c=1.0,tol=0.000001,it0=100,max_iter_backtracking=1000,
#         compute_gram=FALSE,lin_admm=FALSE,admm=FALSE,intercept=FALSE,
#         resetflow=FALSE,regul="",loss="",verbose=FALSE,pos=FALSE,clever=FALSE,
#         log=FALSE,ista=FALSE,subgrad=FALSE,logName="",is_inner_weights=FALSE,
#         inner_weights=c(0.),size_group=1,sqrt_step=TRUE,transpose=FALSE,
#         linesearch_mode=0)
#
# Description:
#   spams.fistaTree solves sparse regularized problems.
#   X is a design matrix of size m x p
#   X=[x^1,...,x^n]', where the x_i's are the rows of X
#   Y=[y^1,...,y^n] is a matrix of size m x n
#   It implements the algorithms FISTA, ISTA and subgradient descent for solving

```

```

#
#     min_W loss(W) + lambda1 psi(W)
#
#     The function psi are those used by spams.proximalTree (see documentation)
#     for the loss functions, see the documentation of spams.fistaFlat
#
#     This function can also handle intercepts (last row of W is not regularized),
#     and/or non-negativity constraints on W and sparse matrices X
#
# Inputs:
#     Y: double dense m x n matrix
#     X: double dense or sparse m x p matrix
#     W0: double dense p x n matrix or p x Nn matrix (for multi-logistic loss)
#         initial guess
#     tree: named list (see documentation of spams.proximalTree)
#     return_optim_info:
#         if true the function will return 2 matrices in a list.
#     loss: (choice of loss, see above)
#     regul: (choice of regularization, see function spams.proximalFlat)
#     lambda1: (regularization parameter)
#     lambda2: (optional, regularization parameter, 0 by default)
#     lambda3: (optional, regularization parameter, 0 by default)
#     verbose: (optional, verbosity level, false by default)
#     pos: (optional, adds positivity constraints on the
#         coefficients, false by default)
#     transpose: (optional, transpose the matrix in the regularization function)
#     size_group: (optional, for regularization functions assuming a group
#         structure)
#     numThreads: (optional, number of threads for exploiting
#         multi-core / multi-cpus. By default, it takes the value -1,
#         which automatically selects all the available CPUs/cores).
#     max_it: (optional, maximum number of iterations, 100 by default)
#     it0: (optional, frequency for computing duality gap, every 10 iterations by default)
#     tol: (optional, tolerance for stopping criterion, which is a relative duality gap
#         if it is available, or a relative change of parameters).
#     gamma: (optional, multiplier for increasing the parameter L in fista, 1.5 by default)
#     L0: (optional, initial parameter L in fista, 0.1 by default, should be small enough)
#     fixed_step: (deactivate the line search for L in fista and use L0 instead)
#     compute_gram: (optional, pre-compute X^TX, false by default).
#     intercept: (optional, do not regularize last row of W, false by default).
#     ista: (optional, use ista instead of fista, false by default).
#     subgrad: (optional, if not ista, use subgradient descent instead of fista, false by
#     default).
#     a:
#     b: (optional, if subgrad, the gradient step is a/(t+b)
#     also similar options as spams.proximalTree
#
#     the function also implements the ADMM algorithm via an option admm=true. It is not
#     documented
#     and you need to look at the source code to use it.
#     delta:    undocumented; modify at your own risks!
#     c:        undocumented; modify at your own risks!
#     max_iter_backtracking:    undocumented; modify at your own risks!
#     lin_admm:    undocumented; modify at your own risks!

```



```

#   admm:    undocumented; modify at your own risks!
#   resetflow:  undocumented; modify at your own risks!
#   clever:   undocumented; modify at your own risks!
#   log:     undocumented; modify at your own risks!
#   logName:  undocumented; modify at your own risks!
#   is_inner_weights:  undocumented; modify at your own risks!
#   inner_weights:  undocumented; modify at your own risks!
#   sqrt_step:  undocumented; modify at your own risks!
#
# Output:
#   W: double dense p x n matrix or p x Nn matrix (for multi-logistic loss)
#   optim: optional, double dense 4 x n matrix.
#         first row: values of the objective functions.
#         third row: values of the relative duality gap (if available)
#         fourth row: number of iterations
#   optim_info: vector of size 4, containing information of the optimization.
#         W <- spams.fistaTree(Y,X,W0,tree,return_optim_info = FALSE,...)
#         v <- spams.fistaTree(Y,X,W0,tree,return_optim_info = TRUE,...)
#         W <- v[[1]]
#         optim_info <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   Valid values for the regularization parameter (regul) are:
#   "l0", "l1", "l2", "l1f", "l2-not-squared", "elastic-net", "fused-lasso",
#   "group-lasso-l2", "group-lasso-l1f", "sparse-group-lasso-l2",
#   "sparse-group-lasso-l1f", "l1l2", "l1l1f", "l1l2+l1", "l1l1f+l1",
#   "tree-l0", "tree-l2", "tree-l1f", "graph", "graph-ridge", "graph-l2",
#   "multi-task-tree", "multi-task-graph", "l1l1f-row-column", "trace-norm",
#   "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
m = 100;n = 10
X = matrix(rnorm(m * n),nrow = m,ncol = n,byrow = FALSE)
X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = spams.normalize(X)
Y = matrix(rnorm(m),nrow = m,ncol = 1,byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
own_variables = as.vector(c(0,0,3,5,6,6,8,9),mode= 'integer')
N_own_variables = as.vector(c(0,3,2,1,0,2,1,1),mode= 'integer')
eta_g = as.vector(c(1,1,1,2,2,2,2.5,2.5),mode = 'double')
groups = matrix(as.vector(c(0, 0, 0, 0, 0, 0, 0, 0,
  1, 0, 0, 0, 0, 0, 0, 0, 0,
  0, 1, 0, 0, 0, 0, 0, 0, 0,
  0, 1, 0, 0, 0, 0, 0, 0, 0,

```

```

1, 0, 0, 0, 0, 0, 0, 0, 0,
0, 0, 0, 0, 1, 0, 0, 0,
0, 0, 0, 0, 1, 0, 0, 0,
0, 0, 0, 0, 0, 0, 1, 0),mode='logical'),ncol = 8,byrow = T)
groups = as(groups,'CsparseMatrix')
tree = list('eta_g'= eta_g,'groups' = groups,'own_variables' = own_variables,
           'N_own_variables' = N_own_variables)
.printf('\nVarious regression experiments\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e-5, intercept = FALSE,pos = FALSE
,compute_gram = TRUE, loss = 'square',regul = 'tree-l2')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])
###
.printf('\nFISTA + Regression tree-linf\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e-5, intercept = FALSE,pos = FALSE
,compute_gram = TRUE, loss = 'square',regul = 'tree-linf')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
[1],optim_info[3],optim_info[4])
###
# works also with non tree-structured regularization. tree is ignored
.printf('\nFISTA + Regression Fused-Lasso\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = TRUE, loss = 'square',regul = 'fused-
lasso')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])
###
.printf('\nISTA + Regression tree-l0\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = TRUE, loss = 'square',regul = 'tree-l0'))
,n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])
###
.printf('\nFISTA + Regression tree-l2 with intercept\n')
x1 = cbind(X,matrix(1,nrow = nrow(X),ncol = 1))
W01 = rbind(W0,matrix(0,nrow = 1, ncol = ncol(W0)))
res = Xtest1('spams',quote(spams.fistaTree(Y,x1,W01,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e
-5, intercept = TRUE,pos = FALSE,compute_gram = TRUE, loss = 'square',regul = 'tree-l2')),n
= 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])

```

```

# Classification

.printf('\nOne classification experiment')
Y = matrix(2. * as.double(rnorm(100 * ncol(Y)) > 0.) - 1.,nrow = 100,ncol = ncol(Y),byrow =
FALSE)
.printf('\nFISTA + Logistic + tree-linf\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,LO = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = TRUE, loss = 'logistic',regul = 'tree-linf
')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
[1],optim_info[3],optim_info[4])
###
# can be used of course with other regularization functions, intercept,...

# Multi-Class classification
Y = ceiling(5 * matrix(runif(100 * ncol(Y),0,1),nrow = 100,ncol = ncol(Y),byrow = FALSE)) - 1
.printf('\nFISTA + Multi-Class Logistic + tree-l2\n')
nclasses = max(Y) + 1
W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,LO = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = TRUE, loss = 'multi-logistic',regul = '
tree-l2')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, number of iterations: %f\n",optim_info[1],optim_info[4])
# can be used of course with other regularization functions, intercept,...

# Multi-Task regression
Y = matrix(rnorm(100 * 100),nrow = 100,ncol = 100,byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
.printf('\nFISTA + Regression multi-task-tree\n')
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = TRUE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,LO = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = FALSE, loss = 'square',regul = 'multi-task
-tree')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
[1],optim_info[3],optim_info[4])

# Multi-Task Classification
.printf('\nFISTA + Logistic + multi-task-tree\n')
Y = matrix(rnorm(100 * ncol(Y)),nrow = 100,ncol = ncol(Y),byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = TRUE,
lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,LO = 0.1, tol = 1e
-5, intercept = FALSE,pos = FALSE,compute_gram = FALSE, loss = 'logistic',regul = 'multi-
task-tree')),n = 1)
W = res[[1]]

```

```

optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

# Multi-Class + Multi-Task Regularization
.printf('\nFISTA + Multi-Class Logistic +multi-task-tree\n')
Y = ceiling(5 * matrix(runif(100 * ncol(Y),0,1),nrow = 100,ncol = ncol(Y),byrow = FALSE)) - 1
nclasses = max(Y) + 1
W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaTree(Y,X,W0,tree,TRUE,numThreads = -1,verbose = FALSE,
  lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,L0 = 0.1, tol = 1e
  -5, intercept = FALSE,pos = FALSE,compute_gram = FALSE, loss = 'multi-logistic',regul = '
  multi-task-tree')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])
# can be used of course with other regularization functions, intercept,...

.printf('\nFISTA + Multi-Class Logistic +multi-task-tree + sparse matrix\n')
nclasses = max(Y) + 1
W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
res = Xtest1('spams',quote(spams.fistaTree(Y,as(X,'CsparseMatrix'),W0,tree,TRUE,numThreads =
  -1,verbose = FALSE,lambda1 = 0.001,lambda2 = 0.001,lambda3 = 0.001, it0 = 10, max_it = 200,
  L0 = 0.1, tol = 1e-5, intercept = FALSE,pos = FALSE,compute_gram = FALSE, loss = 'multi-
  logistic',regul = 'multi-task-tree')),n = 1)
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, number of iterations: %f\n",optim_info
  [1],optim_info[3],optim_info[4])

```

5.10 Function spams.fistaGraph

Given a matrix $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^p]^T$ in $\mathbb{R}^{m \times p}$, and a matrix $\mathbf{Y} = [\mathbf{y}^1, \dots, \mathbf{y}^n]$, it solves the optimization problems presented in the previous section, with the same regularization functions as spams.proximalGraph. see usage details below:

```

#
# Name: spams.fistaGraph
#
# Usage: spams.fistaGraph(Y,X,W0,graph,return_optim_info = FALSE,numThreads = -1,max_it = 1000,L0
  = 1.0,
#         fixed_step=FALSE,gamma=1.5,lambda1=1.0,delta=1.0,lambda2=0.,lambda3=0.,
#         a=1.0,b=0.,c=1.0,tol=0.000001,it0=100,max_iter_backtracking=1000,
#         compute_gram=FALSE,lin_admm=FALSE,admm=FALSE,intercept=FALSE,
#         resetflow=FALSE,regul="",loss="",verbose=FALSE,pos=FALSE,clever=FALSE,
#         log=FALSE,ista=FALSE,subgrad=FALSE,logName="",is_inner_weights=FALSE,
#         inner_weights=c(0.),size_group=1,sqrt_step=TRUE,transpose=FALSE,
#         linesearch_mode=0)
#
# Description:
#   spams.fistaGraph solves sparse regularized problems.
#   X is a design matrix of size m x p
#   X=[x^1,...,x^n]', where the x_i's are the rows of X
#   Y=[y^1,...,y^n] is a matrix of size m x n

```

```

# It implements the algorithms FISTA, ISTA and subgradient descent.
#
# It implements the algorithms FISTA, ISTA and subgradient descent for solving
#
#   min_W loss(W) + lambda1 psi(W)
#
# The function psi are those used by spams.proximalGraph (see documentation)
# for the loss functions, see the documentation of spams.fistaFlat
#
# This function can also handle intercepts (last row of W is not regularized),
# and/or non-negativity constraints on W.
#
# Inputs:
#   Y: double dense m x n matrix
#   X: double dense or sparse m x p matrix
#   W0: double dense p x n matrix or p x Nn matrix (for multi-logistic loss)
#       initial guess
#   graph: struct (see documentation of spams.proximalGraph)
#   return_optim_info:
#       if true the function will return 2 matrices in a list.
#   loss: (choice of loss, see above)
#   regul: (choice of regularization, see function spams.proximalFlat)
#   lambda1: (regularization parameter)
#   lambda2: (optional, regularization parameter, 0 by default)
#   lambda3: (optional, regularization parameter, 0 by default)
#   verbose: (optional, verbosity level, false by default)
#   pos: (optional, adds positivity constraints on the
#        coefficients, false by default)
#   numThreads: (optional, number of threads for exploiting
#               multi-core / multi-cpus. By default, it takes the value -1,
#               which automatically selects all the available CPUs/cores).
#   max_it: (optional, maximum number of iterations, 100 by default)
#   it0: (optional, frequency for computing duality gap, every 10 iterations by default)
#   tol: (optional, tolerance for stopping criterion, which is a relative duality gap
#        if it is available, or a relative change of parameters).
#   gamma: (optional, multiplier for increasing the parameter L in fista, 1.5 by default)
#   L0: (optional, initial parameter L in fista, 0.1 by default, should be small enough)
#   fixed_step: (deactivate the line search for L in fista and use L0 instead)
#   compute_gram: (optional, pre-compute X^TX, false by default).
#   intercept: (optional, do not regularize last row of W, false by default).
#   ista: (optional, use ista instead of fista, false by default).
#   subgrad: (optional, if not ista, use subgradient descent instead of fista, false by
#            default).
#   a:
#   b: (optional, if subgrad, the gradient step is a/(t+b)
#   also similar options as spams.proximalTree
#
# the function also implements the ADMM algorithm via an option admm=true. It is not
# documented
# and you need to look at the source code to use it.
# delta:   undocumented; modify at your own risks!
# c:       undocumented; modify at your own risks!
# max_iter_backtracking:   undocumented; modify at your own risks!
# lin_admm:   undocumented; modify at your own risks!

```

```

#   admm:    undocumented; modify at your own risks!
#   resetflow:  undocumented; modify at your own risks!
#   clever:    undocumented; modify at your own risks!
#   log:      undocumented; modify at your own risks!
#   logName:   undocumented; modify at your own risks!
#   is_inner_weights:  undocumented; modify at your own risks!
#   inner_weights:  undocumented; modify at your own risks!
#   sqrt_step:    undocumented; modify at your own risks!
#   size_group:   undocumented; modify at your own risks!
#   transpose:   undocumented; modify at your own risks!
#
# Output:
#   W:  double dense p x n matrix or p x Nn matrix (for multi-logistic loss)
#   optim: optional, double dense 4 x n matrix.
#       first row: values of the objective functions.
#       third row: values of the relative duality gap (if available)
#       fourth row: number of iterations
#   optim_info:      vector of size 4, containing information of the optimization.
#       W <- spams.fistaGraph(Y,X,W0,graph,return_optim_info = FALSE,...)
#       v <- spams.fistaGraph(Y,X,W0,graph,return_optim_info = TRUE,...)
#       W <- v[[1]]
#       optim_info <- v[[2]]
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
# Note:
#   Valid values for the regularization parameter (regul) are:
#   "l0", "l1", "l2", "linf", "l2-not-squared", "elastic-net", "fused-lasso",
#   "group-lasso-l2", "group-lasso-linf", "sparse-group-lasso-l2",
#   "sparse-group-lasso-linf", "l1l2", "l1linf", "l1l2+l1", "l1linf+l1",
#   "tree-l0", "tree-l2", "tree-linf", "graph", "graph-ridge", "graph-l2",
#   "multi-task-tree", "multi-task-graph", "l1linf-row-column", "trace-norm",
#   "trace-norm-vec", "rank", "rank-vec", "none"
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
set.seed(0)
num_threads = -1 # all cores (-1 by default)
verbose = FALSE # verbosity, false by default
lambda1 = 0.1 # regularization ter
it0 = 1 # frequency for duality gap computations
max_it = 100 # maximum number of iterations
L0 = 0.1
tol = 1e-5
intercept = FALSE
pos = FALSE
eta_g = as.vector(c(1, 1, 1, 1, 1),mode='double')
groups = as(matrix(as.vector(c(0, 0, 0, 1, 0,
0, 0, 0, 0, 0,
0, 0, 0, 0, 0,

```

```

0, 0, 0, 0, 0,
0, 0, 1, 0, 0),mode='logical'),ncol = 5,byrow = T),'CsparseMatrix')
groups_var = as(matrix(as.vector(c(1, 0, 0, 0, 0,
1, 0, 0, 0, 0,
1, 0, 0, 0, 0,
1, 1, 0, 0, 0,
0, 1, 0, 1, 0,
0, 1, 0, 1, 0,
0, 1, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 0, 0, 1,
0, 0, 1, 0, 0),mode='logical'),ncol = 5,byrow = T),'CsparseMatrix')
graph = list('eta_g' = eta_g, 'groups' = groups, 'groups_var' = groups_var)
verbose = TRUE
X = matrix(rnorm(1000),nrow = 100,ncol = 10,byrow = FALSE)
X = X - matrix(rep(colMeans(X),nrow(X)),nrow(X),ncol(X),byrow = T)
X = spams.normalize(X)
Y = matrix(rnorm(100),nrow = 100,ncol = 1,byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
# Regression experiments
# 100 regression problems with the same design matrix X.
.printf('\nVarious regression experiments\n')
compute_gram = TRUE
#
.printf('\nFISTA + Regression graph\n')
loss = 'square'
regul = 'graph'
tic = proc.time()
res = spams.fistaGraph(
  Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
  lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
  intercept = intercept,pos = pos,compute_gram = compute_gram,
  loss = loss,regul = regul)
tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
  optim_info[1],optim_info[3],t,optim_info[4])
#
.printf('\nADMM + Regression graph\n')
admm = TRUE
lin_admm = TRUE
c = 1
delta = 1
tic = proc.time()
res = spams.fistaGraph(
  Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
  lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
  intercept = intercept,pos = pos,compute_gram = compute_gram,
  loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,

```



```

    c = c,delta = delta)
tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
        optim_info[1],optim_info[3],t,optim_info[4])
#
admm = FALSE
max_it = 5
it0 = 1
tic = proc.time()
res = spams.fistaGraph(
    Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
    lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
    intercept = intercept,pos = pos,compute_gram = compute_gram,
    loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
    c = c,delta = delta)
tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time:%f, number of iterations: %f\n",
        optim_info[1],optim_info[3],t,optim_info[4])
#
# works also with non graph-structured regularization. graph is ignored
.printf('\nFISTA + Regression Fused-Lasso\n')
regul = 'fused-lasso'
lambda2 = 0.01
lambda3 = 0.01
tic = proc.time()
res = spams.fistaGraph(
    Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
    lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
    intercept = intercept,pos = pos,compute_gram = compute_gram,
    loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
    c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)
tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, time: %f, number of iterations: %f\n",optim_info[1],t,optim_info[4])
#
.printf("\nFISTA + Regression graph with intercept\n")
regul = 'graph'
intercept = TRUE
tic = proc.time()
res = spams.fistaGraph(
    Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
    lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
    intercept = intercept,pos = pos,compute_gram = compute_gram,
    loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
    c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

```



```

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
    optim_info[1],optim_info[3],t,optim_info[4])
intercept = FALSE
# Classification
.printf('\nOne classification experiment\n')
Y = matrix(2. * as.double(rnorm(100 * ncol(Y)) > 0.) - 1.,nrow = 100,ncol = ncol(Y),byrow =
    FALSE)
.printf('\nFISTA + Logistic + graph-linf\n')
loss = 'logistic'
regul = 'graph'
lambda1 = 0.01
tic = proc.time()
res = spams.fistaGraph(
    Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
    lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
    intercept = intercept,pos = pos,compute_gram = compute_gram,
    loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
    c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
    optim_info[1],optim_info[3],t,optim_info[4])
#
# can be used of course with other regularization functions, intercept,...

# Multi-Class classification
Y = ceiling(5 * matrix(runif(100 * ncol(Y),0,1),nrow = 100,ncol = ncol(Y),byrow = FALSE)) - 1
loss = 'multi-logistic'
regul = 'graph'
.printf('\nFISTA + Multi-Class Logistic + graph\n')
nclasses = max(Y) + 1
W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
tic = proc.time()
res = spams.fistaGraph(
    Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
    lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
    intercept = intercept,pos = pos,compute_gram = compute_gram,
    loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
    c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
    optim_info[1],optim_info[3],t,optim_info[4])
#

```

```

# can be used of course with other regularization functions, intercept,...

# Multi-Task regression
Y = matrix(rnorm(100 * ncol(Y)),nrow = 100,ncol = ncol(Y),byrow = FALSE)
Y = Y - matrix(rep(colMeans(Y),nrow(Y)),nrow(Y),ncol(Y),byrow = T)
Y = spams.normalize(Y)
W0 = matrix(c(0),nrow = ncol(X), ncol = ncol(Y))
compute_gram = FALSE
verbose = TRUE
loss = 'square'
.printf('\nFISTA + Regression multi-task-graph\n')
regul = 'multi-task-graph'
lambda2 = 0.01
tic = proc.time()
res = spams.fistaGraph(
  Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
  lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
  intercept = intercept,pos = pos,compute_gram = compute_gram,
  loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
  c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
  optim_info[1],optim_info[3],t,optim_info[4])
#
# Multi-Task Classification
.printf('\nFISTA + Logistic + multi-task-graph\n')
regul = 'multi-task-graph'
lambda2 = 0.01
loss = 'logistic'
Y = matrix(rnorm(100 * ncol(Y)),nrow = 100,ncol = ncol(Y),byrow = FALSE)
tic = proc.time()
res = spams.fistaGraph(
  Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
  lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
  intercept = intercept,pos = pos,compute_gram = compute_gram,
  loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
  c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
  optim_info[1],optim_info[3],t,optim_info[4])
#
# Multi-Class + Multi-Task Regularization
verbose = FALSE
.printf('\nFISTA + Multi-Class Logistic +multi-task-graph\n')
Y = ceiling(5 * matrix(runif(100 * ncol(Y),0,1),nrow = 100,ncol = ncol(Y),byrow = FALSE)) - 1
nclasses = max(Y) + 1

```

```

W0 = matrix(0,nrow = ncol(X),nclasses * ncol(Y),byrow = FALSE)
loss = 'multi-logistic'
regul = 'multi-task-graph'
tic = proc.time()
res = spams.fistaGraph(
  Y,X,W0,graph,TRUE,numThreads = num_threads,verbose = verbose,
  lambda1 = lambda1,it0 = it0,max_it = max_it,L0 = L0,tol = tol,
  intercept = intercept,pos = pos,compute_gram = compute_gram,
  loss = loss,regul = regul,admm = admm,lin_admm = lin_admm,
  c = c,lambda2 = lambda2,lambda3 = lambda3,delta = delta)

tac = proc.time()
t = (tac - tic)[['elapsed']]
W = res[[1]]
optim_info = res[[2]]
.printf("mean loss: %f, mean relative duality_gap: %f, time: %f, number of iterations: %f\n",
  optim_info[1],optim_info[3],t,optim_info[4])
#
# can be used of course with other regularization functions, intercept,...

```

5.11 Function spams.fistaPathCoding

Similarly, the toolbox handles the penalties of [24] with the following function

```

#
# The R function is not yet implemented.
#

```

The following piece of code illustrates how to use this function.

5.12 Function solverPoisson

The following problem is addressed here

$$\min_{\mathbf{w} \in \mathbb{R}_+^p} \sum_{j=1}^p \mathbf{x}_j^\top \mathbf{w} + \delta - y_j \log(\mathbf{x}_j^\top \mathbf{w} + \delta) + \lambda \psi(\mathbf{w}).$$

```

#
# The R function is not yet implemented.
#

```

The following piece of code illustrates how to use this function.

5.13 Function mexIncrementalProx

This implements the incremental solver MISO [25].

```

#
# The R function is not yet implemented.
#

```

The following piece of code illustrates how to use this function.

5.14 Function mexStochasticProx

This implements the stochastic proximal gradient solver [26].

```
#  
# The R function is not yet implemented.  
#
```

The following piece of code illustrates how to use this function.

6 A few Functions for manipulating images

This functions are not well documented yet

6.1 Function mexExtractPatches

```
#  
# The R function is not yet implemented.  
#
```

6.2 Function mexCombinePatches

```
#  
# The R function is not yet implemented.  
#
```

6.3 Function mexConvFistaFlat

```
#  
# The R function is not yet implemented.  
#
```

7 Miscellaneous Functions

7.1 Function spams.conjGrad

Implementation of a conjugate gradient for solving a linear system $\mathbf{Ax} = \mathbf{b}$ when \mathbf{A} is positive definite. In some cases, it is faster than the Matlab function `pcg`, especially when the library uses the Intel Math Kernel Library.

```
#  
# Name: spams.conjGrad  
#  
# Usage: spams.conjGrad(A,b,x0 = NULL,tol = 1e-10,itermax = NULL)  
#  
# Description:  
#   Conjugate gradient algorithm, sometimes faster than the  
#   equivalent R function solve. In order to solve Ax=b;  
#  
# Inputs:  
#   A: double square n x n matrix. HAS TO BE POSITIVE DEFINITE  
#   b: double vector of length n.  
#   x0: double vector of length n. (optional) initial guess.  
#   tol: (optional) tolerance.  
#   itermax: (optional) maximum number of iterations.  
#
```

```

# Output:
#       x: double vector of length n.
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
A = matrix(rnorm(500 * 5000),nrow = 5000,ncol = 500)
A = t(A) %% A
b = as.vector(rep(1,ncol(A)))
x0 = b
tol = 1e-4
itermax = floor(0.5 * length(b))
CGtest <- function(txt,expr) {
  tic = proc.time()
  for (i in 1:20)
    y = eval(expr)
  tac = proc.time()
  cat(sprintf("  Time (%s): %f\n",txt,(tac - tic)[['user.self']]))
  x1 = abs(b - (A %% y))
  cat(sprintf("Mean error on b : %f\n\n",sum(x1) / length(b)))
}
y2 = CGtest("R",quote(solve(A,b)))
y1 = CGtest("spams",quote(spams.conjGrad(A,b,x0,tol,itermax)))

```

7.2 Function spams.bayer

Apply a Bayer pattern to an input image

```

#
# Name: spams.conjGrad
#
# Usage: spams.conjGrad(A,b,x0 = NULL,tol = 1e-10,itermax = NULL)
#
# Description:
#   Conjugate gradient algorithm, sometimes faster than the
#   equivalent R function solve. In order to solve Ax=b;
#
# Inputs:
#   A: double square n x n matrix. HAS TO BE POSITIVE DEFINITE
#   b: double vector of length n.
#   x0: double vector of length n. (optional) initial guess.
#   tol: (optional) tolerance.
#   itermax: (optional) maximum number of iterations.
#
# Output:
#   x: double vector of length n.
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)

```

```
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
A = matrix(rnorm(500 * 5000),nrow = 5000,ncol = 500)
A = t(A) %*% A
b = as.vector(rep(1,ncol(A)))
x0 = b
tol = 1e-4
itermax = floor(0.5 * length(b))
CGtest <- function(txt,expr) {
  tic = proc.time()
  for (i in 1:20)
    y = eval(expr)
  tac = proc.time()
  cat(sprintf("  Time (%s): %f\n",txt,(tac - tic)[['user.self']]))
  x1 = abs(b - (A %*% y))
  cat(sprintf("Mean error on b : %f\n\n",sum(x1) / length(b)))
}
y2 = CGtest("R",quote(solve(A,b)))
y1 = CGtest("spams",quote(spams.conjGrad(A,b,x0,tol,itermax)))
```

7.3 Function `spams.calcAAAt`

For an input sparse matrix A , this function returns AA^T . For some reasons, when A has a lot more columns than rows, this function can be much faster than the equivalent Matlab command $X \cdot A'$.

```
#
# Name: spams.calcAAAt
#
# Usage: spams.calcAAAt(A)
#
# Description:
#   Compute efficiently  $AA^T = A \cdot A'$ , when  $A$  is sparse
#   and has a lot more columns than rows. In some cases, it is
#   up to 20 times faster than the equivalent R expression
#    $AA^T = A \cdot A'$ ;
#
# Inputs:
#   A: double sparse m x n matrix
#
# Output:
#   AAAt: double m x m matrix
#
# Authors:
#   Julien MAIRAL, 2009 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
m = 200;n = 200000; d= 0.05
```

```
A = rSpMatrix(m,n,floor(m * n * d))
spams.calcAAAt(A)
```

7.4 Function spams.calcXAt

For an input sparse matrix \mathbf{A} and a matrix \mathbf{X} , this function returns \mathbf{XA}^T . For some reasons, when \mathbf{A} has a lot more columns than rows, this function can be much faster than the equivalent Matlab command $\mathbf{X*A}'$.

```
#
# Name: spams.calcXAt
#
# Usage: spams.calcXAt(X,A)
#
# Description:
#   Compute efficiently  $\mathbf{XA}^T = \mathbf{X*A}'$ , when  $\mathbf{A}$  is sparse and has a
#   lot more columns than rows. In some cases, it is up to 20 times
#   faster than the equivalent R expression;
#
# Inputs:
#   X: double m x n matrix
#   A: double sparse p x n matrix
#
# Output:
#   XAt: double m x p matrix
#
# Authors:
#   Julien MAIRAL, 2009 (spams, matlab interface and documentation)
#   Jean-Paul CHIEZE 2011-2012 (R interface)
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
m = 200;n = 200000; d= 0.05
A = rSpMatrix(m,n,floor(m * n * d))
X = matrix(rnorm(64 * n),nrow = 64,ncol = n)
spams.calcXAt(X,A)
```

7.5 Function spams.calcXY

For two input matrices \mathbf{X} and \mathbf{Y} , this function returns \mathbf{XY} .

```
#
# Name: spams.calcXY
#
# Usage: spams.calcXY(X,Y)
#
# Description:
#   Compute  $\mathbf{Z=XY}$  using the BLAS library used by SPAMS.
#
# Inputs:
#   X: double m x n matrix
#   Y: double n x p matrix
#
# Output:
```

```
#      Z: double m x p matrix
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
X = matrix(rnorm(64 * 200),nrow = 64,ncol = 200,byrow = FALSE)
Y = matrix(rnorm(200 * 20000),nrow = 200,ncol = 20000,byrow = FALSE)
spams.calcXY(X,Y)
```

7.6 Function spams.calcXYt

For two input matrices \mathbf{X} and \mathbf{Y} , this function returns \mathbf{XY}^T .

```
#
# Name: spams.calcXYt
#
# Usage: spams.calcXYt(X,Y)
#
# Description:
#      Compute  $Z=XY'$  using the BLAS library used by SPAMS.
#
# Inputs:
#      X: double m x n matrix
#      Y: double p x n matrix
#
# Output:
#      Z: double m x p matrix
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#
```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```
library(spams)
X = matrix(rnorm(64 * 200),nrow = 64,ncol = 200,byrow = FALSE)
Y = matrix(rnorm(200 * 20000),nrow = 20000,ncol = 200,byrow = FALSE)
spams.calcXYt(X,Y)
```

7.7 Function spams.calcXtY

For two input matrices \mathbf{X} and \mathbf{Y} , this function returns $\mathbf{X}^T\mathbf{Y}$.

```
#
# Name: spams.calcXtY
#
# Usage: spams.calcXtY(X,Y)
#
# Description:
```



```

#   Compute Z=X'Y using the BLAS library used by SPAMS.
#
# Inputs:
#   X:  double n x m matrix
#   Y:  double n x p matrix
#
# Output:
#   Z:  double m x p matrix
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
X = matrix(rnorm(64 * 200),nrow = 200,ncol = 64,byrow = FALSE)
Y = matrix(rnorm(200 * 20000),nrow = 200,ncol = 20000,byrow = FALSE)
spams.calcXtY(X,Y)

```

7.8 Function spams.invSym

For an input symmetric matrices \mathbf{A} in $\mathbb{R}^{n \times n}$, this function returns \mathbf{A}^{-1} .

```

#
# Name: spams.invSym
#
# Usage: spams.invSym(A)
#
# Description:
#   returns the inverse of a symmetric matrix A
#
# Inputs:
#   A:  double n x n matrix
#
# Output:
#   B:  double n x n matrix
#
# Authors:
# Julien MAIRAL, 2009 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
A = matrix(runif(1000 * 1000,0,1),nrow = 1000,ncol = 1000,byrow = FALSE)
A = t(A) %% A
spams.invSym(A)

```

7.9 Function spams.normalize

```

#
# Name: spams.normalize
#
# Usage: spams.normalize(X)
#
# Description:
#     rescale the columns of X so that they have
#     unit l2-norm.
#
# Inputs:
#     X: double m x n matrix
#
# Output:
#     Y: double m x n matrix
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
A = matrix(runif(100 * 1000,0,1),nrow = 100,ncol = 1000)
y2 = spams.normalize(A)

```

7.10 Function spams.sort

```

#
# Name: spams.sort
#
# Usage: spams.sort(X,mode=T)
#
# Description:
#     sort the elements of X using quicksort
#
# Inputs:
#     X: double vector of size n
#     mode: false for decreasing order (true by default)
#
# Output:
#     Y: double vector of size n
#
# Authors:
# Julien MAIRAL, 2010 (spams, matlab interface and documentation)
# Jean-Paul CHIEZE 2011-2012 (R interface)
#

```

The following piece of code illustrates how to use this function. The following piece of code contains usage examples:

```

library(spams)
x = rnorm(2000000,0,1)
spams.sort(x,TRUE)

```

7.11 Function `displayPatches`

Print to the screen a matrix containing as columns image patches.

7.12 Function `spams.countPathsDAG`

This function counts the number of paths in a DAG using dynamic programming.

```
#  
# The R function is not yet implemented.  
#
```

The following piece of code illustrates how to use this function.

7.13 Function `spams.removeCyclesGraph`

One heuristic to remove cycles from a graph.

```
#  
# The R function is not yet implemented.  
#
```

The following piece of code illustrates how to use this function.

7.14 Function `spams.countConnexComponents`

Count the number of connected components of a subgraph from a graph.

```
#  
# The R function is not yet implemented.  
#
```

The following piece of code illustrates how to use this function.

7.15 Function `spams.graphOfGroupStruct`

```
#  
# Name: spams.graphOfGroupStruct  
#  
# Usage: spams.graphOfGroupStruct(gstruct)  
#  
# Description:  
#   converts a group structure into the graph structure  
#   used by spams.proximalGraph, spams.fistaGraph or spams.structTrainDL  
#  
# Inputs:  
#   gstruct: the structure of groups as a list, one element per node  
#   an element is itself a 4 elements list:  
#     nodeid (>= 0), weight (double), array of vars associated to the node,  
#     array of children (nodeis's)  
#  
# Output:  
#   graph: struct (see documentation of spams.proximalGraph)  
#  
# Authors:  
# Jean-Paul CHIEZE, 2012  
#
```

7.16 Function `spams.groupStructOfString`

```
#
# Name: spams.groupStructOfString
#
# Usage: spams.groupStructOfString(s)
#
# Description:
#   decode a multi-line string describing "simply" the structure of groups
#   of variables needed by spams.proximalGraph, spams.proximalTree, spams.fistaGraph,
#   spams.fistaTree and spams.structTrainDL and builds the corresponding group structure.
#   Each line describes a group of variables as a node of a tree.
#   It has up to 4 fields separated by spaces:
#       node-id node-weight [variables-list] -> children-list
#   Let's define Ng = number of groups, and Nv = number of variables.
#   node-id must be in the range (0 - Ng-1), and there must be Ng nodes
#   weight is a float
#   variables-list : a space separated list of integers, maybe empty,
#       but '[' and ']' must be present. Numbers in the range (0 - Nv-1)
#   children-list : a space separated list of node-id's
#       If the list is empty, '->' may be omitted.
#   The data must obey some rules :
#       - A group contains the variables of the corresponding node and of the whole subtree.
#       - Variables attached to a node are those that are not in the subtree.
#       - If the data destination is a Graph, there may be several independant trees,
#         and a variable may appear in several trees.
#   If the destination is a Tree, there must be only one tree, the root node
#   must have id == 0 and each variable must appear only once.
#
# Inputs:
#   s: the multi-lines string
#
# Output:
#   groups: list, one element for each node
#           an element is itself a 4 elements list:
#           nodeid (int >= 0), weight (double), array of vars of the node,
#           array of children (nodeid's)
#
# Authors:
#   Jean-Paul CHIEZE, 2012
#
```

7.17 Function `spams.readGroupStruct`

```
#
# Name: spams.readGroupStruct
#
# Usage: spams.readGroupStruct(file)
#
# Description:
#   reads a text file describing "simply" the structure of groups
#   of variables needed by spams.proximalGraph, spams.proximalTree, spams.fistaGraph,
#   spams.fistaTree and spams.structTrainDL and builds the corresponding group structure.
#   weight is a float
```

```

# variables-list : a space separated list of integers, maybe empty,
# but '[' and ']' must be present. Numbers in the range (0 - Nv-1)
# children-list : a space separated list of node-id's
# If the list is empty, '->' may be omitted.
# The data must obey some rules :
# - A group contains the variables of the corresponding node and of the whole subtree.
# - Variables attached to a node are those that are not in the subtree.
# - If the data destination is a Graph, there may be several independent trees,
# and a variable may appear in several trees.
# If the destination is a Tree, there must be only one tree, the root node
# must have id == 0 and each variable must appear only once.
#
# Inputs:
# file: the file name
#
# Output:
# groups: list, one element for each node
# an element is itself a 4 elements list:
# nodeid (int >= 0), weight (double), array of vars of the node,
# array of children (nodeid's)
#
# Authors:
# Jean-Paul CHIEZE, 2012
#

```

7.18 Function `spams.treeOfGroupStruct`

```

#
# Name: spams.treeOfGroupStruct
#
# Usage: spams.treeOfGroupStruct(gstruct)
#
# Description:
# converts a group structure into the tree structure
# used by spams.proximalTree, spams.fistaTree or spams.structTrainDL
#
# Inputs:
# gstruct: the structure of groups as a list, one element per node
# an element is itself a 4 elements list:
# nodeid (>= 0), weight (double), array of vars associated to the node,
# array of children (nodeid's)
#
# Output:
# permutations: permutation vector that must be applied to the result of the
# programm using the tree. Empty if no permutation is needed.
# tree: named list (see documentation of spams.proximalTree)
# nbvars : number of variables in the tree
# res <- spams.treeOfGroupStruct(gstruct)
# permutations = res[[1]]
# tree = res[[2]]
# nbvars = res[[3]]
#
# Authors:
# Jean-Paul CHIEZE, 2012

```

```
#
```

7.19 Function spams.simpleGroupTree

```
#
# Name: spams.simpleGroupTree
#
# Usage: spams.simpleGroupTree(degrees)
#
# Description:
#   makes a structure representing a tree given the
#   degree of each level.
#
# Inputs:
#   degrees: int vector; degrees(i) is the number of children of each node at level i
#
# Output:
#   group_struct: list, one element for each node
#   an element is itself a 4 elements list :
#   nodeid (int >= 0), weight (double), array of vars attached to the node
#   (here equal to [nodeid]), array of children (nodeid's)
#
# Authors:
# Jean-Paul CHIEZE, 2012
#
```

A Duality Gaps with Fenchel Duality

This section is taken from the appendix D of Julien Mairal's PhD thesis [19]. We are going to use intensively Fenchel Duality (see [2]). Let us consider the problem

$$\min_{\mathbf{w} \in \mathbb{R}^p} [g(\mathbf{w}) \triangleq f(\mathbf{w}) + \lambda\psi(\mathbf{w})], \quad (47)$$

We first notice that for all the formulations we have been interested in, $g(\mathbf{w})$ can be rewritten

$$g(\mathbf{w}) = \tilde{f}(\mathbf{X}^\top \mathbf{w}) + \lambda\psi(\mathbf{w}), \quad (48)$$

where $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]$ are training vectors, and \tilde{f} is an appropriated smooth real-valued function of \mathbb{R}^n , and ψ one of the regularization functions we have introduced.

Given a primal variable \mathbf{w} in \mathbb{R}^p and a dual variable $\boldsymbol{\kappa}$ in \mathbb{R}^n , we obtain using classical Fenchel duality rules [2], that the following quantity is a duality gap for problem (47):

$$\delta(\mathbf{w}, \boldsymbol{\kappa}) \triangleq g(\mathbf{w}) + \tilde{f}^*(\boldsymbol{\kappa}) + \lambda\psi^*(-\mathbf{X}\boldsymbol{\kappa}/\lambda),$$

where \tilde{f}^* and ψ^* are respectively the Fenchel conjugates of \tilde{f} and ψ . Denoting by \mathbf{w}^* the solution of Eq. (47), the duality gap is interesting in the sense that it upperbounds the difference with the optimal value of the function:

$$\delta(\mathbf{w}, \boldsymbol{\kappa}) \geq g(\mathbf{w}) - g(\mathbf{w}^*) \geq 0.$$

Similarly, we will consider pairs of primal-dual variables (\mathbf{W}, \mathbf{K}) when dealing with matrices.

During the optimization, sequences of primal variables \mathbf{w} are available, and one wishes to exploit duality gaps for estimating the difference $g(\mathbf{w}) - g(\mathbf{w}^*)$. This requires the following components:

- being able to efficiently compute \tilde{f}^* and ψ^* .

- being able to obtain a “good” dual variable $\boldsymbol{\kappa}$ given a primal variable \mathbf{w} , such that $\delta(\mathbf{w}, \boldsymbol{\kappa})$ is close to $g(\mathbf{w}) - g(\mathbf{w}^*)$.

We suppose that the first point is satisfied (we will detail these computations for every loss and regularization functions in the sequel), and explain how to choose $\boldsymbol{\kappa}$ in general (details will also be given in the sequel).

Let us first consider the choice that associates with a primal variable \mathbf{w} , the dual variable

$$\boldsymbol{\kappa}(\mathbf{w}) \triangleq \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w}), \quad (49)$$

and let us compute $\delta(\mathbf{w}, \boldsymbol{\kappa}(\mathbf{w}))$. First, easy computations show that for all vectors \mathbf{z} in \mathbb{R}^n , $\tilde{f}^*(\nabla \tilde{f}(\mathbf{z})) = \mathbf{z}^\top \nabla \tilde{f}(\mathbf{z}) - \tilde{f}(\mathbf{z})$, which gives

$$\delta(\mathbf{w}, \boldsymbol{\kappa}(\mathbf{w})) = \tilde{f}(\mathbf{X}^\top \mathbf{w}) + \lambda \psi(\mathbf{w}) + \tilde{f}^*(\nabla \tilde{f}(\mathbf{X}^\top \mathbf{w})) + \lambda \psi^*(-\mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w})/\lambda), \quad (50)$$

$$= \lambda \psi(\mathbf{w}) + \mathbf{w}^\top \mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w}) + \lambda \psi^*(-\mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w})/\lambda). \quad (51)$$

We now use the classical Fenchel-Young inequality (see, Proposition 3.3.4 of [2]) on the function ψ , which gives

$$\delta(\mathbf{w}, \boldsymbol{\kappa}(\mathbf{w})) \geq \mathbf{w}^\top \mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w}) - \mathbf{w}^\top \mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w}) = 0,$$

with equality if and only if $-\mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w})$ belongs to $\partial \psi(\mathbf{w})$. Interestingly, we now that first-order optimality conditions for Eq. (48) gives that $-\mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w}^*) \in \partial \psi(\mathbf{w}^*)$. We have now in hand a non-negative function $\mathbf{w} \mapsto \delta(\mathbf{w}, \boldsymbol{\kappa}(\mathbf{w}))$ of \mathbf{w} , that upperbounds $g(\mathbf{w}) - g(\mathbf{w}^*)$ and satisfying $\delta(\mathbf{w}^*, \boldsymbol{\kappa}(\mathbf{w}^*)) = 0$.

This is however not a sufficient property to make it a good measure of the quality of the optimization, and further work is required, that will be dependent on \tilde{f} and ψ . We have indeed proven that $\delta(\mathbf{w}^*, \boldsymbol{\kappa}(\mathbf{w}^*))$ is always 0. However, for \mathbf{w} different than \mathbf{w}^* , $\delta(\mathbf{w}^*, \boldsymbol{\kappa}(\mathbf{w}^*))$ can be infinite, making it a non-informative duality-gap. The reasons for this can be one of the following:

- The term $\psi^*(-\mathbf{X} \nabla \tilde{f}(\mathbf{X}^\top \mathbf{w})/\lambda)$ might have an infinite value.
- Intercepts make the problem more complicated. One can write the formulation with an intercept by adding a row to \mathbf{X} filled with the value 1, add one dimension to the vector \mathbf{w} , and consider a regularization function ψ that does regularize the last entry of \mathbf{w} . This further complexifies the computation of ψ^* and its definition, as shown in the next section.

Let us now detail how we proceed to solve these problems, but first without considering the intercept. The analysis is similar when working with matrices \mathbf{W} instead of vectors \mathbf{w} .

A.0.1 Duality Gaps without Intercepts

Let us show how to compute the Fenchel conjugate of the functions we have introduced. We now present the Fenchel conjugate of the loss functions \tilde{f} .

- **The square loss**

$$\begin{aligned} \tilde{f}(\mathbf{z}) &= \frac{1}{2n} \|\mathbf{y} - \mathbf{z}\|_2^2, \\ \tilde{f}^*(\boldsymbol{\kappa}) &= \frac{n}{2} \|\boldsymbol{\kappa}\|_2^2 + \boldsymbol{\kappa}^\top \mathbf{y}. \end{aligned}$$

- **The logistic loss**

$$\begin{aligned} \tilde{f}(\mathbf{z}) &= \frac{1}{n} \sum_{i=1}^n \log(1 + e^{-y_i \mathbf{z}_i}) \\ \tilde{f}^*(\boldsymbol{\kappa}) &= \begin{cases} +\infty & \text{if } \exists i \in [1; n] \text{ s.t. } y_i \boldsymbol{\kappa}_i \notin [-1, 0], \\ \sum_{i=1}^n (1 + y_i \boldsymbol{\kappa}_i) \log(1 + y_i \boldsymbol{\kappa}_i) - y_i \boldsymbol{\kappa}_i \log(-y_i \boldsymbol{\kappa}_i) & \text{otherwise.} \end{cases} \end{aligned}$$

- **The multiclass logistic loss (or softmax).** The primal variable is now a matrix \mathbf{Z} , in $\mathbb{R}^{n \times r}$, which represents the product $\mathbf{X}^\top \mathbf{W}$. We denote by \mathbf{K} the dual variable in $\mathbb{R}^{n \times r}$.

$$\begin{aligned} \tilde{f}(\mathbf{Z}) &= \frac{1}{n} \sum_{i=1}^n \log \left(\sum_{j=1}^r e^{\mathbf{Z}_{ij} - \mathbf{Z}_{iy_i}} \right) \\ \tilde{f}^*(\mathbf{K}) &= \begin{cases} +\infty & \text{if } \exists i \in [1; n] \text{ s.t. } \{\mathbf{K}_{ij} < 0 \text{ and } j \neq y_i\} \text{ or } \mathbf{K}_{iy_i} < -1, \\ \sum_{i=1}^n \left[\sum_{j \neq y_i} \mathbf{K}_{ij} \log(\mathbf{K}_{ij}) + (1 + \mathbf{K}_{iy_i}) \log(1 + \mathbf{K}_{iy_i}) \right]. \end{cases} \end{aligned}$$

Our first remark is that the choice Eq. (49) ensures that $\tilde{f}(\boldsymbol{\kappa})$ is not infinite.

As for the regularization function, except for the Tikhonov regularization which is self-conjugate (it is equal to its Fenchel conjugate), we have considered functions that are norms. There exists therefore a norm $\|\cdot\|$ such that $\psi(\mathbf{w}) = \|\mathbf{w}\|$, and we denote by $\|\cdot\|_*$ its dual-norm. In such a case, the Fenchel conjugate of ψ for a vector $\boldsymbol{\gamma}$ in \mathbb{R}^p takes the form

$$\psi^*(\boldsymbol{\gamma}) = \begin{cases} 0 & \text{if } \|\boldsymbol{\gamma}\|_* \leq 1, \\ +\infty & \text{otherwise.} \end{cases}$$

It turns out that for almost all the norms we have presented, there exists (i) either a closed form for the dual-norm or (ii) there exists an efficient algorithm evaluating it. The only one which does not conform to this statement is the tree-structured sum of ℓ_2 -norms, for which we do not know how to evaluate it efficiently.

One can now slightly modify the definition of $\boldsymbol{\kappa}$ to ensure that $\psi^*(-\mathbf{X}\boldsymbol{\kappa}/\lambda) \neq +\infty$. A natural choice is

$$\boldsymbol{\kappa}(\mathbf{w}) \triangleq \min\left(1, \frac{\lambda}{\|\mathbf{X}\nabla\tilde{f}(\mathbf{X}^\top\mathbf{w})\|_*}\right)\nabla\tilde{f}(\mathbf{X}^\top\mathbf{w}),$$

which is the one we have implemented. With this new choice, it is easy to see that for all vectors \mathbf{w} in \mathbb{R}^p , we still have $\tilde{f}^*(\boldsymbol{\kappa}) \neq +\infty$, and finally, we also have $\delta(\mathbf{w}, \boldsymbol{\kappa}(\mathbf{w})) < +\infty$ and $\delta(\mathbf{w}^*, \boldsymbol{\kappa}(\mathbf{w}^*)) = 0$, making it potentially a good duality gap.

A.0.2 Duality Gaps with Intercepts

Even though adding an intercept does seem a simple modification to the original problem, it induces difficulties for finding good dual variables.

We recall that having an intercept is equivalent to having a problem of the type (48), by adding a row to \mathbf{X} filled with the value 1, adding one dimension to the vector \mathbf{w} (or one row for matrices \mathbf{W}), and by using a regularization function that does not depend on the last entry of \mathbf{w} (or the last row of \mathbf{W}).

Suppose that we are considering a problem of type (48) of dimension $p+1$, but we are using a regularization function $\tilde{\psi} : \mathbb{R}^{p+1} \rightarrow \mathbb{R}$, such that for a vector \mathbf{w} in \mathbb{R}^{p+1} , $\tilde{\psi}(\mathbf{w}) \triangleq \psi(\mathbf{w}_{[1:p]})$, where $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$ is one of the regularization function we have introduced. Then, considering a primal variable \mathbf{w} , a dual variable $\boldsymbol{\kappa}$, and writing $\boldsymbol{\gamma} \triangleq -\mathbf{X}\boldsymbol{\kappa}/\lambda$, we are interested in computing

$$\tilde{\psi}^*(\boldsymbol{\gamma}) = \begin{cases} +\infty & \text{if } \gamma_{p+1} \neq 0 \\ \psi^*(\boldsymbol{\gamma}_{[1:p]}) & \text{otherwise,} \end{cases}$$

which means that in order the duality gap not to be infinite, one needs in addition to ensure that γ_{p+1} be zero. Since the last row of \mathbf{X} is filled with ones, this writes down $\sum_{i=1}^{p+1} \boldsymbol{\kappa}_i = 0$. For the formulation with matrices \mathbf{W} and \mathbf{K} , the constraint is similar but for every column of \mathbf{K} .

Let us now detail how we proceed for every loss function to find a “good” dual variable $\boldsymbol{\kappa}$ satisfying this additional constraint, given a primal variable \mathbf{w} in \mathbb{R}^{p+1} , we first define the auxiliary function

$$\boldsymbol{\kappa}'(\mathbf{w}) \triangleq \nabla\tilde{f}(\mathbf{X}^\top\mathbf{w}),$$

(which becomes $\mathbf{K}'(\mathbf{W}) \triangleq \nabla\tilde{f}(\mathbf{X}^\top\mathbf{W})$ for matrices), and then define another auxiliary function $\boldsymbol{\kappa}''(\mathbf{w})$ as follows, to take into account the additional constraint $\sum_{i=1}^{p+1} \boldsymbol{\kappa}_i = 0$.

- **For the square loss**, we define another auxiliary function:

$$\boldsymbol{\kappa}''(\mathbf{w}) \triangleq \boldsymbol{\kappa}'(\mathbf{w}) - \frac{1}{n}\mathbf{1}_{p+1}^\top\boldsymbol{\kappa}'(\mathbf{w})\mathbf{1}_{p+1}$$

where $\mathbf{1}_{p+1}$ is a vector of size $p+1$ filled with ones. This step, ensures that $\sum_{i=1}^{p+1} \boldsymbol{\kappa}''(\mathbf{w})_i = 0$.

- **For the logistic loss**, the situation is slightly more complicated since additional constraints are involved in the definition of \tilde{f}^* .

$$\boldsymbol{\kappa}''(\mathbf{w}) \triangleq \arg \min_{\boldsymbol{\kappa} \in \mathbb{R}^n} \|\boldsymbol{\kappa} - \boldsymbol{\kappa}'(\mathbf{w})\|_2^2 \quad \text{s.t.} \quad \sum_{i=1}^n \boldsymbol{\kappa}_i = 0 \quad \text{and} \quad \forall i \in [1; n], \boldsymbol{\kappa}_i \in [-1, 0].$$

This problem can be solved in linear-time [3] using a similar algorithm as for the projection onto the ℓ_1 -ball, since it is an instance of a *quadratic knapsack problem*.

- **For the multi-class logistic loss**, we proceed in a similar way, for every column \mathbf{K}^j of \mathbf{K} , $j \in [1; r]$:

$$\mathbf{K}''^j(\mathbf{w}) \triangleq \arg \min_{\boldsymbol{\kappa} \in \mathbb{R}^n} \|\mathbf{K}^j - \boldsymbol{\kappa}'(\mathbf{w})\|_2^2 \quad \text{s.t.} \quad \sum_{i=1}^n \kappa_i = 0 \quad \text{and} \\ \forall i \in [1; n], \{\kappa_i \geq 0 \text{ if } j \neq \mathbf{y}_i\} \text{ and } \{\kappa_i \geq -1 \text{ if } \mathbf{y}_i = j\}.$$

When the function ψ is the Tykhonov regularization function, we end the process by setting $\boldsymbol{\kappa}(\mathbf{w}) = \boldsymbol{\kappa}''(\mathbf{w})$. When it is a norm, we choose, as before for taking into account the constraint $\|\mathbf{X}\boldsymbol{\kappa}\|_* \leq \lambda$,

$$\boldsymbol{\kappa}(\mathbf{w}) \triangleq \min \left(1, \frac{\lambda}{\|\mathbf{X}\boldsymbol{\kappa}''(\mathbf{w})\|_*} \right) \boldsymbol{\kappa}''(\mathbf{w}),$$

with a similar formulation for matrices \mathbf{W} and \mathbf{K} .

Even though finding dual variables while taking into account the intercept requires quite a lot of engineering, notably implementing a quadratic knapsack solver, it can be done efficiently.

References

- [1] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [2] J. M. Borwein and A. S. Lewis. *Convex analysis and nonlinear optimization: Theory and examples*. Springer, 2006.
- [3] P. Brucker. An $O(n)$ algorithm for quadratic knapsack problems. 3:163–166, 1984.
- [4] E. J. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted l1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2008.
- [5] B. V. Cherkassky and A. V. Goldberg. On implementing the push-relabel method for the maximum flow problem. *Algorithmica*, 19(4):390–410, 1997.
- [6] S. F. Cotter, J. Adler, B. Rao, and K. Kreutz-Delgado. Forward sequential algorithms for best basis selection. In *IEEE Proceedings of Vision Image and Signal Processing*, pages 235–244, 1999.
- [7] A. Cutler and L. Breiman. Archetypal analysis. *Technometrics*, 36(4):338–347, 1994.
- [8] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2008.
- [9] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of statistics*, 32(2):407–499, 2004.
- [10] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Annals of statistics*, 1(2):302–332, 2007.
- [11] J. Friedman, T. Hastie, and R. Tibshirani. A note on the group lasso and a sparse group lasso. Technical report, Preprint arXiv:1001.0736, 2010.
- [12] W. J. Fu. Penalized regressions: The bridge versus the Lasso. *Journal of computational and graphical statistics*, 7:397–416, 1998.
- [13] A. V. Goldberg and R. E. Tarjan. A new approach to the maximum flow problem. In *Proc. of ACM Symposium on Theory of Computing*, pages 136–146, 1986.
- [14] P. O. Hoyer. Non-negative sparse coding. In *Proc. IEEE Workshop on Neural Networks for Signal Processing*, 2002.
- [15] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for sparse hierarchical dictionary learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2010.

- [16] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12:2297–2334, 2011.
- [17] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems*, 2001.
- [18] N. Maculan and J. R. G. Galdino de Paula. A linear-time median-finding algorithm for projecting a vector on the simplex of \mathbb{R}^n . *Operations research letters*, 8(4):219–222, 1989.
- [19] J. Mairal. *Sparse coding for machine learning, image processing and computer vision*. PhD thesis, Ecole Normale Supérieure, Cachan, 2010.
- [20] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2009.
- [21] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11:19–60, 2010.
- [22] J. Mairal, R. Jenatton, G. Obozinski, and F. Bach. Network flow algorithms for structured sparsity. In *Advances in Neural Information Processing Systems*, 2010.
- [23] J. Mairal, R. Jenatton, G. Obozinski, and F. Bach. Convex and network flow optimization for structured sparsity. *Journal of Machine Learning Research*, 12:2649–2689, 2011.
- [24] J. Mairal and B. Yu. Supervised feature selection in graphs with path coding penalties and network flows. *Journal of Machine Learning Research*, 2013.
- [25] Julien Mairal. Optimization with first-order surrogate functions. In *International Conference on Machine Learning (ICML)*, 2013.
- [26] Julien Mairal. Stochastic majorization-minimization algorithms for large-scale optimization. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- [27] S. Mallat and Z. Zhang. Matching pursuit in a time-frequency dictionary. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- [28] N. Meinshausen and P. Bühlmann. Stability selection. Technical report. ArXiv:0809.2932.
- [29] G. Obozinski, B. Taskar, and M.I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, pages 1–22.
- [30] M. R. Osborne, B. Presnell, and B. A. Turlach. On the Lasso and its dual. *Journal of Computational and Graphical Statistics*, 9(2):319–37, 2000.
- [31] P. Sprechmann, I. Ramirez, G. Sapiro, and Y. C. Eldar. Collaborative hierarchical sparse modeling. Technical report, 2010. Preprint arXiv:1003.0400v1.
- [32] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society Series B*, 67(1):91–108, 2005.
- [33] J. A. Tropp. Algorithms for simultaneous sparse approximation. part ii: Convex relaxation. *Signal Processing, special issue "Sparse approximations in signal and image processing"*, 86:589–602, April 2006.
- [34] J. A. Tropp, A. C. Gilbert, and M. J. Strauss. Algorithms for simultaneous sparse approximation. part i: Greedy pursuit. *Signal Processing, special issue "sparse approximations in signal and image processing"*, 86:572–588, April 2006.
- [35] S. Weisberg. *Applied Linear Regression*. Wiley, New York, 1980.
- [36] T. T. Wu and K. Lange. Coordinate descent algorithms for Lasso penalized regression. *Annals of Applied Statistics*, 2(1):224–244, 2008.

- [37] J. Mairal Y. Chen and Z. Harchaoui. Fast and robust archetypal analysis for representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [38] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B*, 68:49–67, 2006.
- [39] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B*, 67(2):301–320, 2005.