

Tutorial on usage of “hsdar“-package

Hanna Meyer, Lukas W. Lehnert

June 18, 2015

Contents

1	Introduction	3
1.1	Sample data	3
1.2	How to start	3
2	Handling of speclibs	4
2.1	Structure	4
2.2	Creating Speclibs	5
2.2.1	Speclibs from matrices	5
2.2.2	Speclibs from raster files	7
3	Plotting Speclibs	9
4	Simulating spectra with PROSPECT and PROSAIL	11
5	Basic data manipulation tools	16
5.1	Subsets of spectra	16
5.2	Filter	17
5.3	Mask	19
5.4	Calculations of derivations	21
5.5	Resampling of bands to various satellite sensors	23
6	Continuum removal	24
6.1	Manually adapting continuum lines	26
6.2	Extracting absorption features	30
7	Basic data analysis tools	32
7.1	T-Test	32
7.2	Regressions	33
8	Calculating spectral indices	35
8.1	Vegetation indices	35
8.2	Red edge parameters	36
8.3	Normalised ratio indices	37
8.4	Comparing distributions of NRI	39
9	Analysing relations between NRI and environmental variables	40
9.1	Correlations	40
9.2	Linear models	43
9.3	Generalized linear models	46
9.4	Plot NRI models	46
10	Linear spectral unmixing	48

1 Introduction

This tutorial introduces techniques for creating, handling, manipulating, analyzing and simulating hyperspectral data using the `hsdar` package. Though we'll try our best to keep the examples as easy as possible, we assume that you are already familiar with the R software.

In most cases this tutorial is not built successively so that you can get in at your chapter of interest. However, you should briefly run over this introduction to see how to use this tutorial.

The tutorial focuses on the usage of `hsdar` for the calculation of several common methods in hyperspectral data manipulation and analysis. Despite some less common methods, we won't explain in detail what the methods do and in which cases they are useful or not. Please read the `hsdar` help files and references for more information about your methods of interest.

`hsdar` is still experimental. If you have any questions, suggestions or concerns don't hesitate to contact the authors. For some applications of `hsdar` see [Lehnert et al. \(2013, 2014, 2015\)](#); [Meyer et al. \(2013\)](#).

1.1 Sample data

Almost all of the exercises in this tutorial use one single sample dataset - "spectral_data" - which is included in the `hsdar` package. This dataset was created in scope of a field campaign in Tibet in 2012 aiming to analyze spatial patterns of pasture degradation on the plateau. Two study areas on the plateau are included in this dataset: "Kailash" which is located in the drier western of Tibet and "Namco" which is a more humid site in central Tibet.

In the first line, the dataset contains hyperspectral reflectance measurements which were taken with a field spectrometer from a height of approx. 1.50 m covering the integrated spectrum of a circle of approx. 50 cm on the ground. Thus, reflectance values result from different fractions of vegetation, soil, stones etc. Furthermore, the dataset contains environmental variables of the measured circles like vegetation coverage, chlorophyll content, soil moisture etc.

1.2 How to start

To work with the tutorial, first install the `hsdar` package and load the library as well as the sample data:

```
> #install.packages("hsdar")
> library(hsdar)
> data(spectral_data) #Load the data used in the tutorial
```

If you need help, see

```
> help(hsdar)
```

2 Handling of speclibs

In this chapter a "Speclib" which is the main class of `hsdar` will be presented. Almost all functions of `hsdar` require that your Spectra are stored in a Speclib. To learn how to handle Speclibs, we will first have a look on the structure of the sample data. Afterwards it will be shown how to create own Speclibs and how to read and write them.

2.1 Structure

Hyperspectral data as well as further information related to these data are stored in a class called "Speclib". To understand the structure of a Speclib, have a look on the sample Speclib "spectral_data"

```
> spectral_data #See how Speclibs are printed
```

```
Summary of Speclib
```

```
Summary of spectra
```

```
-----  
Total number of spectra : 81  
Number of bands : 1401  
Width of bands : 0.5
```

The printed information of a Speclib contain the number of spectra, the number of spectral bands and the width of the bands. However, there are more information stored in the Speclib. Have a look on the structure of "spectral_data" to see all its information:

```
> str(spectral_data)
```

```
Formal class 'Speclib' [package "hsdar"] with 13 slots  
..@ spectra          : num [1:81, 1:1401] 9.08 8.63 9.32 9.59 8.54 ...  
..@ attr(*, "dimnames")=List of 2  
.. ..$ : NULL  
.. ..$ : NULL  
..@ wavelength      : int [1:1401] 305 306 307 308 309 310 311 312 313 314 ...  
..@ attributes      : 'data.frame':      81 obs. of  10 variables:  
.. ..$ Site         : chr [1:81] "Namco" "Namco" "Namco" "Namco" ...  
.. ..$ ID           : Factor w/ 81 levels "NC103","Nc119",...: 2 3 4 1 29 30 31 16 17 18 ...  
.. ..$ PV           : num [1:81] 67.5 52.5 32.5 65 72.5 75 45 77.5 60 77.5 ...  
.. ..$ Bare.soil    : num [1:81] 30.2 34.7 49.9 20.3 16.3 ...  
.. ..$ Open.turf    : num [1:81] 0.805 5.904 0 5.27 0.5 ...  
.. ..$ Stones       : num [1:81] 0.3 1.1 6.8 1.2 2.4 0 3.25 3 16.5 0.5 ...  
.. ..$ NPV          : num [1:81] NA NA NA NA NA NA NA NA NA ...  
.. ..$ Soil.moisture : num [1:81] 19.8 21.3 22.3 16.9 23.8 24.3 20.3 18.9 19.8 21.3 ...  
.. ..$ Chlorophyll   : num [1:81] NA NA NA 49.6 45 ...  
.. ..$ Target.species: num [1:81] NA NA NA 1 1 1 NA NA 1 1 ...  
..@ fwhm            : num 0.5  
..@ continuousdata  : logi TRUE  
..@ wlunit          : chr "nm"  
..@ xlabel          : chr "Wavelength"  
..@ ylabel          : chr "Reflectance"  
..@ ID              : chr [1:81] "Nc119" "Nc123" "Nc124" "NC103" ...  
..@ wavelength.is.range: logi FALSE  
..@ transformation  : chr(0)  
..@ usagehistory     : chr(0)  
..@ rastermeta      : list()
```

Only considering the most important components (slots), the `Speclib` contains spectra, wavelengths, different attributes (optional) and some metadata (optional). The spectra are stored in a matrix with the spectral bands organized in columns and the different samples (or pixels) organized in rows. The vector "wavelength" indicates the corresponding wavelength of each band. The "attributes" `data.frame` offers a possibility to store further information related to the spectra, in this case, these are "location", "ID" and several environmental variables like e.g., the soil moisture. Thus, the wavelength contains the metadata of the column and the attributes contain the information of the rows of the spectra-matrix. Finally, some metadata are given like "reflectance" as the type of the spectra or "nm" as the unit of the wavelength.

2.2 Creating Speclibs

Now, we will explain how to create your own `Speclib`s. We will do this the way that we split the sample `Speclib` back into its components and then show how to bring the components together into a new `Speclib`.

2.2.1 Speclibs from matrices

`Speclib`s can be created in different ways. To build a `Speclib` you need at least spectra and the corresponding wavelength values. The easiest way is to prepare a matrix of your spectra. This matrix must be organized in the way that each row represents one sample and each column represents a spectral band. To go on with the example introduced above, we transform the spectra of "spectral_data" back into a matrix:

```
> spectra <- spectra(spectral_data)
```

See what happened:

```
> str(spectra)

num [1:81, 1:1401] 9.08 8.63 9.32 9.59 8.54 ...
- attr(*, "dimnames")=List of 2
 ..$ : NULL
 ..$ : NULL
```

This is how the input matrix must look like: We have a matrix with each row representing one spectrum and each column representing one channel. Further, we need a vector indicating which wavelength corresponds to each column. Therefore we will extract the wavelength from "spectral_data":

```
> wavelength <- wavelength(spectral_data)
```

Now both components needed to create a `Speclib` are available: spectra and the corresponding wavelength. Now you can build a new `Speclib` from them:

```
> newSpecLib <- specLib(spectra, wavelength)
```

Having a look at the structure showing that you have re-created "spectral_data":

```
> str(newSpecLib)

Formal class 'SpecLib' [package "hsdar"] with 13 slots
 ..@ spectra      : num [1:81, 1:1401] 9.08 8.63 9.32 9.59 8.54 ...
 .. ..- attr(*, "dimnames")=List of 2
 .. .. ..$ : NULL
 .. .. ..$ : NULL
 ..@ wavelength   : int [1:1401] 305 306 307 308 309 310 311 312 313 314 ...
 ..@ attributes   : 'data.frame': 0 obs. of 0 variables
 ..@ fwhm         : num 1
 ..@ continuousdata : logi TRUE
 ..@ wlunit       : chr "nm"
 ..@ xlabel       : chr "Wavelength"
 ..@ ylabel       : chr "Reflectance"
 ..@ ID          : chr(0)
 ..@ wavelength.is.range: logi FALSE
 ..@ transformation : chr(0)
 ..@ usagehistory  : chr(0)
 ..@ rastermeta    : list()
```

However, it would be nice to have an ID for each spectrum:

```
> ids <- idSpecLib(spectral_data) #extract ID from "spectral_data"
> idSpecLib(newSpecLib) <- as.character(ids) #...and assign them to the new SpecLib
```

Still the attributes are missing in the new SpecLib. Those can be handled with the function "attribute": If you extract the attributes from "spectral_data" using "attribute", you get a data.frame containing the values for each attribute at each site:

```
> attributes <- attribute(spectral_data)
> head(attributes)
```

	Site	ID	PV	Bare.soil	Open.turf	Stones	NPV	Soil.moisture
1	Namco	Nc119	67.5	30.1975	0.805	0.3	NA	19.8
2	Namco	Nc123	52.5	34.7480	5.904	1.1	NA	21.3
3	Namco	Nc124	32.5	49.8500	0.000	6.8	NA	22.3
4	Namco	NC103	65.0	20.2650	5.270	1.2	NA	16.9
5	Namco	NC92	72.5	16.3000	0.500	2.4	NA	23.8
6	Namco	NC94	75.0	14.5000	0.500	0.0	NA	24.3
	Chlorophyll Target.species							
1		NA		NA				
2		NA		NA				
3		NA		NA				
4		49.63333		1				
5		45.04167		1				
6		45.23333		1				

You can now use this data.frame to complete your new Speclib:

```
> attribute(newSpeclib) <- attributes
```

Finally you have a Speclib which is well comparable to the exemplary Speclib:

```
> str(newSpeclib)

Formal class 'Speclib' [package "hsdar"] with 13 slots
 ..@ spectra      : num [1:81, 1:1401] 9.08 8.63 9.32 9.59 8.54 ...
 ..@ attr(*, "dimnames")=List of 2
 .. ..$ : NULL
 .. ..$ : NULL
 ..@ wavelength   : int [1:1401] 305 306 307 308 309 310 311 312 313 314 ...
 ..@ attributes   : 'data.frame':      81 obs. of  10 variables:
 .. ..$ Site      : chr [1:81] "Namco" "Namco" "Namco" "Namco" ...
 .. ..$ ID        : Factor w/ 81 levels "NC103","Nc119",...: 2 3 4 1 29 30 31 16 17 18 ...
 .. ..$ PV        : num [1:81] 67.5 52.5 32.5 65 72.5 75 45 77.5 60 77.5 ...
 .. ..$ Bare.soil  : num [1:81] 30.2 34.7 49.9 20.3 16.3 ...
 .. ..$ Open.turf  : num [1:81] 0.805 5.904 0 5.27 0.5 ...
 .. ..$ Stones    : num [1:81] 0.3 1.1 6.8 1.2 2.4 0 3.25 3 16.5 0.5 ...
 .. ..$ NPV       : num [1:81] NA NA NA NA NA NA NA NA NA ...
 .. ..$ Soil.moisture : num [1:81] 19.8 21.3 22.3 16.9 23.8 24.3 20.3 18.9 19.8 21.3 ...
 .. ..$ Chlorophyll : num [1:81] NA NA NA 49.6 45 ...
 .. ..$ Target.species: num [1:81] NA NA NA 1 1 1 NA NA 1 1 ...
 ..@ fwhm         : num 1
 ..@ continuousdata : logi TRUE
 ..@ wlunit       : chr "nm"
 ..@ xlabel       : chr "Wavelength"
 ..@ ylabel       : chr "Reflectance"
 ..@ ID          : chr [1:81] "Nc119" "Nc123" "Nc124" "NC103" ...
 ..@ wavelength.is.range: logi FALSE
 ..@ transformation : chr(0)
 ..@ usagehistory  : chr(0)
 ..@ rastermeta    : list()
```

2.2.2 Speclibs from raster files

The investigation of data taken with a hyperspectral camera or by hyperspectral satellite sensors (e.g., Hyperion) brings along that a huge number of spectra must be analyzed. Unless you are working on a large cluster, R won't be able to store all of the data in RAM in this case. A workaround is, to analyze each row of the image separately. This is possible as far as you are using techniques which do not take the neighboring spectra into account. This is the case for almost all functions in this tutorial above except the correlation and linear regression techniques. A very good starting point for the row-wise analysis of large raster files is given in the tutorials of the "raster"-package the **hsdar**-package is depending on. Nevertheless, there are two specific difficulties if hyperspectral data should be analyzed:

1. The wavelength information must be stored along the spectra
2. Most of the function in the **hsdar**-package require that data is transferred to functions as Speclib.

Thus, a small extension of the raster-classes is provided by the **hsdar**-package: the class **HyperSpecRaster**. This class is more or less a **RasterBrick**-object with wavelength information.

In the following example which is taken from the help page of **HyperSpecRaster**, we will first create a small hyperspectral raster file using **PROSAIL** (for explanation of **PROSAIL** see section 4):

```

> ## Create raster file using PROSAIL
> ## Run PROSAIL
> parameter <- data.frame(N = c(rep.int(seq(0.5, 1.4, 0.1), 6)),
+                           LAI = c(rep.int(0.5, 10), rep.int(1, 10),
+                                   rep.int(1.5, 10), rep.int(2, 10),
+                                   rep.int(2.5, 10), rep.int(3, 10)))
> spectra <- PROSAIL(parameterList = parameter)
> ## Create SpatialPixelsDataFrame and fill data with spectra from
> ## PROSAIL
> rows <- round(nspectra(spectra)/10, 0)
> cols <- ceiling(nspectra(spectra)/rows)
> grd <- SpatialGrid(GridTopology(cellcentre.offset = c(1,1,1),
+                                   cellsize = c(1,1,1),
+                                   cells.dim = c(cols, rows, 1)))
> x <- SpatialPixelsDataFrame(grd,
+                             data = as.data.frame(spectra(spectra)))
> ## Write data to example file (example_in.tif) in workingdirectory
> writeGDAL(x, fname = "example_in.tif", drivename = "GTiff")

```

Once the file is created, we will read it back into R and create an object of class `HyperSpecRaster` from it:

```

> infile <- "example_in.tif"
> wavelength <- spectra$wavelength
> ra <- HyperSpecRaster(infile, wavelength)
> tr <- blockSize(ra)

```

Note that we haven't read the values of the file into memory, yet. This is now performed in a small loop over all rows: Let's assume that we want to calculate all available vegetation indices from the hyperspectral image. Thus, we will read each row into memory, calculate the vegetation indices from the subset of pixels in the memory and store the output in a new file:

```

> outfile <- "example_result.tif"
> n_veg <- as.numeric(length(vegindex()))
> res <- writeStart(ra, outfile, overwrite = TRUE, nl = n_veg)
> for (i in 1:tr$n)
+ {
+   v <- getValuesBlock(ra, row=tr$row[i], nrows=tr$nrows[i])
+   mask(v) <- c(1350, 1450)
+   v <- as.matrix(vegindex(v, index=vegindex()))
+   res <- writeValues(res, v, tr$row[i])
+ }
> res <- writeStop(res)

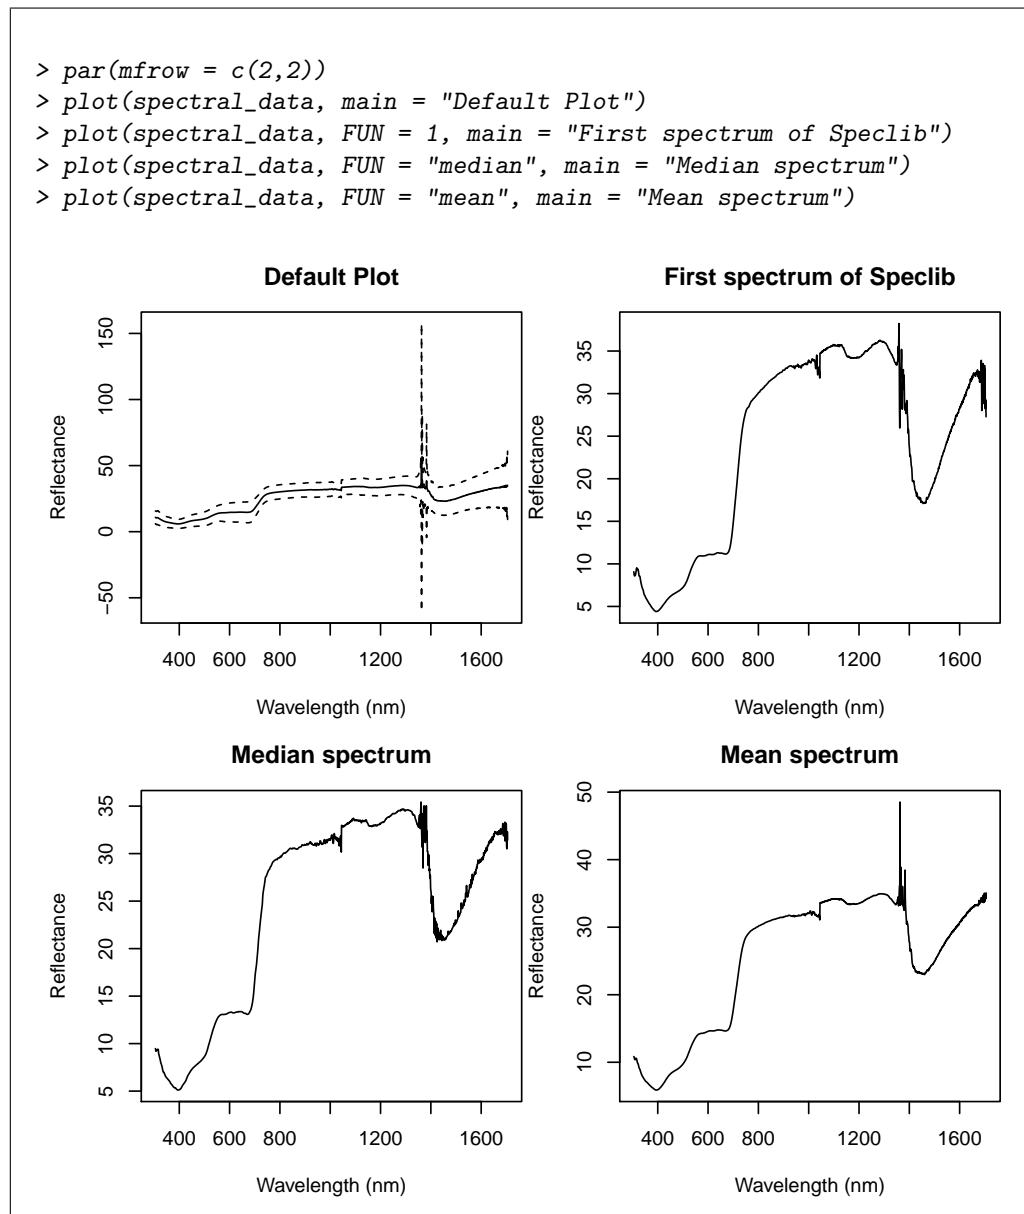
```

Note that `hsdar` is automatically transferring the data to an object of class `Speclib` during each step in the loop (so, `v` is a `Speclib`).

Depending on the number of columns in your image and on the amount of memory of your computer the loop may significantly speed up if you read multiple rows per iteration step. See the tutorial in the raster package mentioned above for further examples and information.


3 Plotting Speclibs

Speclibs can easily be plotted using the `plot.speclib` function. The default way is to plot mean values (solid line) of all spectra in the Speclib and the standard deviations within bands. If the data are continuous the standard deviations are plotted as dashed lines otherwise error bars will indicate standard deviations. You can also plot single spectra by adapting the `FUN` parameter to the ID of the spectra to be plotted. Also, you can use a function as `FUN` parameter like e.g. the median or mean spectrum. See some examples below:




There are some more parameters which might be interesting to plot the data: The parameter "new" allows you to plot more than one spectrum in one plot:

```
> plot(spectral_data, FUN = 1, col = "red")  
> plot(spectral_data, FUN = 2, col = "blue", new = FALSE)  
> plot(spectral_data, FUN = 3, col = "orange", new = FALSE)
```



With the "subset" parameter you can "zoom" into the plot by setting the minimum and maximum wavelength to plot:

```
> plot(spectral_data, subset = c(600,800),  
+      legend = list(x = "topleft", bty = "n")) #add legend
```




Beside these specific arguments to plot spectra, any arguments known from the default plot function can be used.

4 Simulating spectra with PROSPECT and PROSAIL

PROSPECT is a widely used leaf reflectance (or transmittance) model which simulates reflectance values between 400 and 2500 nm. For a detailed description of PROSPECT see [Jacquemoud and Baret \(1990\)](#). PROSPECT requires a set of parameters describing structure and chemical composition of leaves. In `hsdar` there are default values for each parameter. However, these default values were included with the intention to provide an easy access to the model and should be used with care in any scientific approach! But now, let's jump into the simulation of reflectance values:

```
> ## Simulate first spectrum with lower chlorophyll content
> spectrum1 <- PROSPECT(N = 1.3, Cab = 30, Car = 10, Cbrown = 0,
+                       Cw = 0.01, Cm = 0.01)
> ## Simulate second spectrum with higher chlorophyll content
> spectrum2 <- PROSPECT(N = 1.3, Cab = 60, Car = 10, Cbrown = 0,
+                       Cw = 0.01, Cm = 0.01)
> ## Plot results:
> plot(spectrum1, col = "darkorange4", ylim = c(0,0.5),
+      subset = c(400, 800))
> plot(spectrum2, col = "darkgreen", new = FALSE)
```



In addition to PROSPECT, PROSAIL simulates the reflectance of the canopy of vegetation. Thus, the number of parameters is considerably larger and includes the geometry of the plants and the viewing and illumination geometry. General information about PROSAIL may be found in [Jacquemoud et al. \(2009\)](#). In the following example, we will use another way to specify the parameters which is available in PROSPECT and PROSAIL: the `parameterList`. Let's assume we want to test the effect of the illumination geometry (the solar zenith angle) on reference values. It would be possible to simulate many different spectra with different parameter settings. However, we would then have a confusing number of speclibs. An easier way is the `parameterList` (which is a `data.frame` in R):

```
> ## Defining parameter
> parameter <- data.frame(tts = seq(15, 85, 0.5))
> head(parameter)
  tts
1 15.0
2 15.5
3 16.0
4 16.5
5 17.0
6 17.5

> ## Perform simulation (all other parameters are set to default
> ## values)
```

```

> spectra <- PROSAIL(parameterList = parameter)
> spectra

Summary of Speclib

Summary of spectra
-----
Total number of spectra : 141
Number of bands : 2101
Width of bands : 1

> ## Let's see the attributes
> summary(spectra$attributes)

      N      Cab      Car      Cbrown      Cw
Min.   :1.5   Min.   :40   Min.   :8   Min.   :0   Min.   :0.01
1st Qu.:1.5   1st Qu.:40   1st Qu.:8   1st Qu.:0   1st Qu.:0.01
Median :1.5   Median :40   Median :8   Median :0   Median :0.01
Mean   :1.5   Mean   :40   Mean   :8   Mean   :0   Mean   :0.01
3rd Qu.:1.5   3rd Qu.:40   3rd Qu.:8   3rd Qu.:0   3rd Qu.:0.01
Max.   :1.5   Max.   :40   Max.   :8   Max.   :0   Max.   :0.01

      Cm      psoil      LAI      TypeLidf      lidfa
Min.   :0.009   Min.   :0   Min.   :1   Min.   :1   Min.   : -0.35
1st Qu.:0.009   1st Qu.:0   1st Qu.:1   1st Qu.:1   1st Qu.: -0.35
Median :0.009   Median :0   Median :1   Median :1   Median : -0.35
Mean   :0.009   Mean   :0   Mean   :1   Mean   :1   Mean   : -0.35
3rd Qu.:0.009   3rd Qu.:0   3rd Qu.:1   3rd Qu.:1   3rd Qu.: -0.35
Max.   :0.009   Max.   :0   Max.   :1   Max.   :1   Max.   : -0.35

      lidfb      hspot      tts      tto
Min.   : -0.15   Min.   :0.01   Min.   :15.0   Min.   :10
1st Qu.: -0.15   1st Qu.:0.01   1st Qu.:32.5   1st Qu.:10
Median : -0.15   Median :0.01   Median :50.0   Median :10
Mean   : -0.15   Mean   :0.01   Mean   :50.0   Mean   :10
3rd Qu.: -0.15   3rd Qu.:0.01   3rd Qu.:67.5   3rd Qu.:10
Max.   : -0.15   Max.   :0.01   Max.   :85.0   Max.   :10

      psi
Min.   :0
1st Qu.:0
Median :0
Mean   :0
3rd Qu.:0
Max.   :0

```

We can visualize the effect of the solar zenith angle simply plotting the spectra in different colours:

```


> colours <- colorRamp(c("darkorange4", "yellow"))
> plot(spectra, FUN = 1, ylim = c(0, 0.3),
+      col = rgb(colours(spectra$attributes$tts[1]/85),

```

```

+             maxColorValue = 255))
> for (i in 2:nspectra(spectra))
+   plot(spectra, FUN = i, new = FALSE,
+        col = rgb(colours(spectra$attributes$tts[i]/85),
+             maxColorValue = 255))

```



Of course it is also possible to test the effect of two parameters on the reflectance values. In the following example we will plot the reflectance values of canopies with three different LAI values. Within each LAI class we vary the leaf angle distribution:

```

> ## Defining parameter
> parameter <- data.frame(LAI = rep.int(c(1,2,3),5),
+                          TypeLidf = 1,
+                          lidfa = c(rep.int(1,3), rep.int(-1,3),
+                                    rep.int(0,6), rep.int(-0.35,3)),
+                          lidfb = c(rep.int(0,6), rep.int(-1,3),
+                                    rep.int(1,3), rep.int(-0.15,3)))
> parameter

```

	LAI	TypeLidf	lidfa	lidfb
1	1	1	1.00	0.00
2	2	1	1.00	0.00
3	3	1	1.00	0.00
4	1	1	-1.00	0.00

```

5    2    1 -1.00  0.00
6    3    1 -1.00  0.00
7    1    1  0.00 -1.00
8    2    1  0.00 -1.00
9    3    1  0.00 -1.00
10   1    1  0.00  1.00
11   2    1  0.00  1.00
12   3    1  0.00  1.00
13   1    1 -0.35 -0.15
14   2    1 -0.35 -0.15
15   3    1 -0.35 -0.15

```

```

> ## Perform simulation
> spectra <- PROSAIL(parameterList = parameter)
> spectra

```

Summary of SpecLib

Summary of spectra

```


-----
Total number of spectra : 15
Number of bands : 2101
Width of bands : 1

```

```

> ## Plot result:
> ## Colour indicates LAI
> ## Line style indicates LIDF type
> colours <- c("darkblue", "red", "darkgreen")
> LIDF_type <- as.factor(c(rep.int("Planophile", 3),
+                           rep.int("Erectophile", 3),
+                           rep.int("Plagiophile", 3),
+                           rep.int("Extremophile", 3),
+                           rep.int("Spherical", 3)))
> plot(spectra, FUN = 1, ylim = c(0, 0.5),
+       col = colours[spectra$attributes$LAI[1]],
+       lty = which(levels(LIDF_type) == LIDF_type[1]))
> for (i in 2:nspectra(spectra))
+   plot(spectra, FUN = i, new= FALSE,
+         col = colours[spectra$attributes$LAI[i]],
+         lty = which(levels(LIDF_type) == LIDF_type[i]))
> legend("topright",
+       legend = c(paste("LAI =", c(1:3)), "", levels(LIDF_type)),
+       col = c(colours,
+               rep.int("black", 1 + length(levels(LIDF_type)))),
+       lty = c(rep.int(1, 3), 0, 1:length(levels(LIDF_type))))

```



5 Basic data manipulation tools

The `speclib`-class provides several routines for data manipulation which will be described in the following.

5.1 Subsets of spectra

Subsets of `Speclibs` can be built using the `subset` function. This function separates the spectra according to a condition. Usually the conditions are derived from attributes stored in the `Speclib` like study site or vegetation type. For example you could split "spectral_data" to get a subset for the Kailash site and a subset for the Namco site:

```
> ## Return names of attributes data
> names(spectral_data$attributes)

[1] "Site"          "ID"            "PV"
[4] "Bare.soil"     "Open.turf"     "Stones"
[7] "NPV"          "Soil.moisture" "Chlorophyll"
[10] "Target.species"


> ## Devide into both locations using to the attribute "SITE"
> sp_namco <- subset(spectral_data, Site == "Namco")
> sp_kailash <- subset(spectral_data, Site == "Kailash")
```



```

> #
> #Plot results:
> #
> plot(sp_namco, FUN = "mean", col = "darkgreen", ylim = c(0,50))
> plot(sp_kailash, FUN = "mean", col = "darkred", new = FALSE)

```



As you can see the Speclib is split into two Speclibs, one containing all spectra which had the site-attribute "Namco" and one containing all spectra which had the site-attribute "Kailash".


5.2 Filter

Having a look at one spectrum in detail you might wish to smooth the spectrum somehow.

```

> plot(spectral_data, FUN = 1, subset = c(1200,1300)) #raw spectrum

```



`hsdar` implements several methods to smooth spectra. These are Savitzky-Golay, Spline, locally weighted scatterplot smoothing (Lowess) and Mean-filter. `smoothSpecLib` needs a `SpecLib` as input and the method to be used. Depending on the method there are further parameters to be set. Have a look on the `hsdar` help to find more information on these additional parameters.

```
> #
> #Filter SpecLib:
> #
> sgolay <- smoothSpecLib(spectral_data, method = "sgolay", n = 25)
> lowess <- smoothSpecLib(spectral_data, method = "lowess", f = .01)
> meanflt <- smoothSpecLib(spectral_data, method = "mean", p = 5)
> spline <- smoothSpecLib(spectral_data, method = "spline",
+                           n = round(nbands(spectral_data)/10, 0))
```

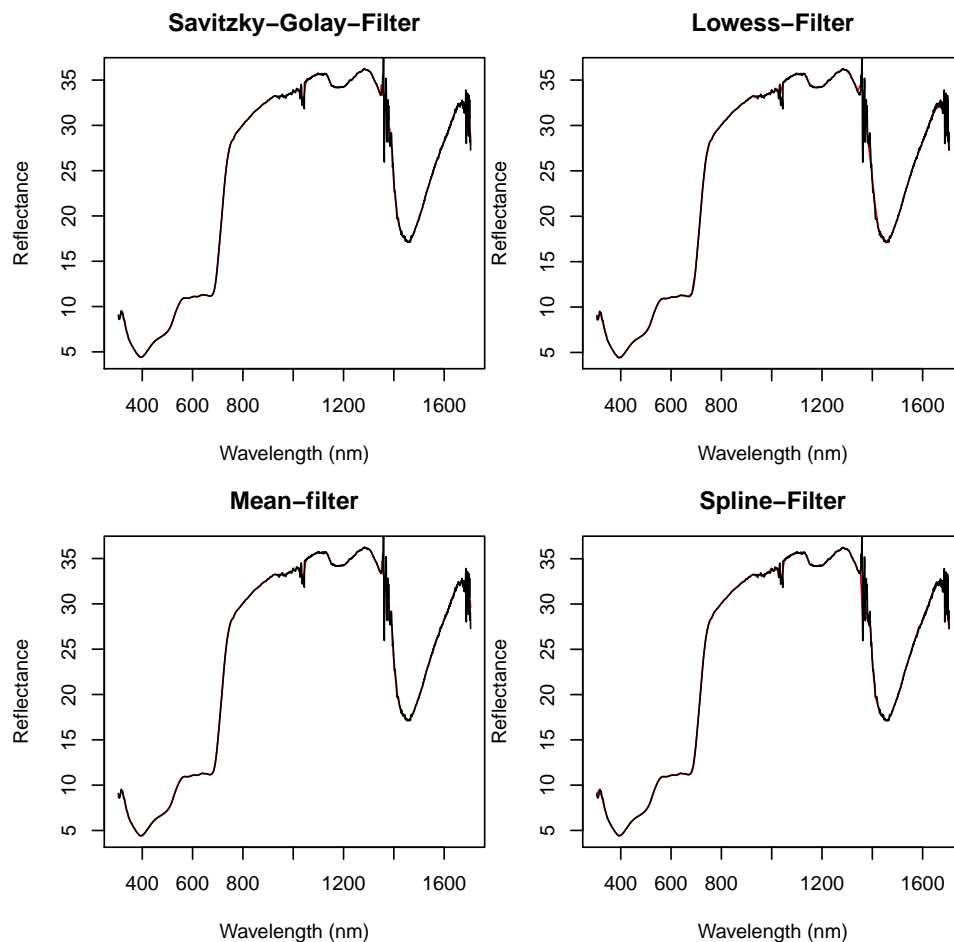
Now plot the results:

```
> par(mfrow = c(2,2))
> plot(sgolay, FUN = 1, subset = c(1200,1300), col = "red",
+ main = "Savitzky-Golay-Filter")
```

```

> plot(spectral_data, FUN = 1, new = FALSE) #raw spectrum
> plot(lowess, FUN = 1, subset = c(1200,1300), col = "red",
+ main = "Lowess-Filter")
> plot(spectral_data, FUN = 1, new = FALSE) #raw spectrum
> plot(meanflt, FUN = 1, subset = c(1200,1300), col = "red",
+ main = "Mean-filter")
> plot(spectral_data, FUN = 1, new = FALSE) #raw spectrum
> plot(spline, FUN = 1, subset = c(1200,1300), col = "red",
+ main = "Spline-Filter")
> plot(spectral_data, FUN = 1, new = FALSE) #raw spectrum

```




5.3 Mask

Usually, there are parts in the spectrum which are associated with errors or which are simply not of interest. `hsdar` allows you to mask these parts so that they don't appear in further analysis any more. In "spectral_data", the areas between 1040 and 1060 nm are errors due to channel crossing of the spectrometer and the wavelengths 1300 to 1450 are affected by water absorption. These areas should be masked in the following. There are several ways of how to enter the lower and upper limits of the wavelengths to be masked. For example you can set these values from a vector which simply consists of a sequence of lower and upper wavelengths. All wavelength between lower and upper wavelength are


then masked. See `?mask` for further options of how to specify these values.

```
> spectral_data_masked <- spectral_data
> mask(spectral_data_masked) <- c(1040,1060,1300,1450)
> #
> #plot results:
> #
> par(mfrow = c(1,2))
> plot(spectral_data, FUN = 1)
> plot(spectral_data_masked, FUN = 1)
```



Beside of masking these wavelength you can also assign "new" values to them by linear interpolation. Note that interpolation is not working if start or end point of the whole spectrum were masked.


```
> spectral_data_interpolated <- interpolate.mask(spectral_data_masked)
> plot(spectral_data_interpolated, FUN = 1)
```



5.4 Calculations of derivations


The derivation of spectra are needed in some analyzing techniques for example to characterize the shape of the red edge. The number of derivation is indicated by the parameter m , thus $m = 1$ returns the first derivation of the spectra.

```
> spectral_data_1deriv <- derivative.speclib(spectral_data, m = 1)
> spectral_data_2deriv <- derivative.speclib(spectral_data, m = 2)
> #
> #plot derivations of the red edge area of 1. spectrum in the Speclib:
> #
> par(mfrow=c(1,2))
> plot(spectral_data_1deriv, FUN = 1, subset = c(600,800),
+ main = "First derivation")
> plot(spectral_data_2deriv, FUN = 1, subset = c(600,800),
+ main = "Second Derivation")
```



However, depending on the objectives it might be favorable to first smooth the spectra before calculating the derivations:

```
> spectral_data_1deriv <- derivative.speclib(smoothSpecLib(
+ spectral_data, method = "sgolay", n = 35), m = 1)
> spectral_data_2deriv <- derivative.speclib(smoothSpecLib(
+ spectral_data, method = "sgolay", n = 35), m = 2)
> #
> #Plot results:
> #
> par(mfrow=c(1,2))
> plot(spectral_data_1deriv, FUN = 1, subset = c(600,800),
+ main = "First derivation")
> plot(spectral_data_2deriv, FUN = 1, subset = c(600,800),
+ main = "Second Derivation")
```



5.5 Resampling of bands to various satellite sensors

`hsdar` allows to resample the `speclib` to the bands of common satellite sensors. The characteristics of (satellite) sensor to integrate spectra can be chosen from a list of already implemented sensors or they can be passed as a `data.frame` with two columns: first column with lower bounds of channels and second column with upper bounds. See which sensors are already implemented:

```
> get.sensor.characteristics(0)
```

	Numerical	abbreviation	Response function implemented
ALI	5		FALSE
EnMAP	11		FALSE
Hyperion	6		FALSE
Landsat4	9		TRUE
Landsat5	4		TRUE
Landsat7	10		TRUE
Landsat8	12		TRUE
MODIS	1		FALSE
Quickbird	7		TRUE
RapidEye	2		TRUE
WorldView2-4	8		TRUE
WorldView2-8	3		TRUE

For some sensors the spectral response functions are available. The spectra can be resampled in three ways. One possibility is the use of the spectral response functions, if available. Otherwise spectra can be resampled assuming a Gaussian distribution (`responsefunction = FALSE`) or the mean value (`responsefunction = NA`) of reflectances between the limiting wavelength of a satellites channel.

```
> ## use spectral response function
> spectral_data_resampled <- spectral.resampling(spectral_data,
+                                                "WorldView2-8")
```

See what changed in the `Speclib`:

```
> spectral_data_resampled
```

Summary of `Speclib`

History of usage

(1) Integrated spectra to `WorldView2-8` channels

Summary of spectra

```


-----
Total number of spectra : 81
Number of bands : 8
Mean width of bands : 66.25

> wavelength(spectral_data_resampled)

[1] 427 478 546 608 659 724 831 908

> #
> #plot results:
> plot(spectral_data_resampled)

```



You can see that the number of bands is now reduced to the 8 WorldView2 bands instead of 1401 hyperspectral channels. The wavelengths are automatically adapted in the `speclib` attribute. Note that standard deviations are now not longer plotted as dashed line by default but as error bars because the data are not longer continuous.

6 Continuum removal

Continuum removal is a commonly used method in hyperspectral remote sensing to normalize spectra and to detect and ensure the comparability of absorption features. The continuum removal transformation is performed by firstly establishing a continuum line/hull which connects the local maxima of the reflectance spectrum. Two kinds of this hull are well established in scientific community: the convex hull (e.g. [Mutanga and Skidmore](#)

(2004a)) and the segmented hull (e.g. [Clark et al. \(1987\)](#)). Both hulls are established by connecting the local maxima, however, the precondition of the convex hull is that the resulting continuum line must be convex whereas considering the segmented hull it might be concave or convex but the algebraic sign of the slope is not allowed to change from the global maximum of the spectrum downwards to the sides. In contrast to a convex hull, the segmented hull is able to identify small absorption features.

Because the continuum removal transformation is sensitive to errors in the spectrum, it's advisable to first mask erroneous parts of the spectrum. Further you should consider to smooth the spectra very slightly to avoid small local maxima which are not associated with reflection maxima. Read section [5.2](#) and [5.3](#) for more details on the preprocessing.

```
> spectral_data_preproc <- smoothSpeclib(spectral_data,
+ method = "sgolay", n = 5)
> mask(spectral_data_preproc) <- c(1040,1060,1300,1450)
```

Then have a look at the transform.speclib function:

```
> str(transform_speclib)

function (data, ..., method = "ch", out = "bd")
```


"data" defines the speclib which is to be transformed. Concerning the "methods" parameter, currently the mentioned **convex hull** ("ch") and the **segmented hull** ("sh") are implemented. The "out" parameter indicates whether the continuum line ("raw"), the continuum removed spectra (**b**and **d**epth, "bd") or the "ratio" will be returned. Have a look on the help page of transform.speclib and the listed literature for details on these methods or for help with interpretation. The output type of "ratio" the "bd" is a speclib, out="raw" returns an object of "clman". "clman" is a class designed to store and handle manual continuum lines.

The following example will show you how to calculate the continuum line (just for visualization) and the band depth using the convex hull as well as the segmented hull approach:

```
> #convex hull:
> ch_cline <- transform_speclib(spectral_data_preproc,
+                               method = "ch", out = "raw")
> ch_bd <- transform_speclib(spectral_data_preproc,
+                             method = "ch", out = "bd")
> #
> #segmented hull:
> #
> sh_cline <- transform_speclib(spectral_data_preproc,
+                               method = "sh", out = "raw")
> sh_bd <- transform_speclib(spectral_data_preproc,
+                             method = "sh", out = "bd")
```

Plot continuum lines and resulting band depths for both methods to see the differences:

```
> #plot results for the first spectrum:
> #
> par(mfrow = c(2,2))
> plot(ch_cline, ispec = 1, numeratepoints = FALSE,
+      main = "Convex hull - Continuum line")
> plot(ch_bd, ispec = 1, main = "Convex hull - Band depth")
> plot(sh_cline, ispec = 1, numeratepoints = FALSE,
+      main = "Segmented hull - Continuum line")
> plot(sh_bd, ispec = 1, main = "Segmented hull - Band depth")
```



6.1 Manually adapting continuum lines


Let's have a look on the segmented hull in more detail and compare it to another spectrum. Therefore, we zoom to the red edge area:

```
> par(mfrow = c(1,2))
> plot(sh_cline, ispec = 1, main = "Continuum line, Spectrum 1",
```

```

+ subset = c(500,800)) #first spectrum
> plot(sh_cline, ispec = 5, main = "Continuum line, Spectrum 5",
+ subset = c(500,800)) #fifth spectrum

```



By the way: Plotting an object of the class "clman" allows you to numerate the local maxima which were used to construct this line.

Comparing the first spectrum and the fifth spectrum it is obvious that in the first there are several small local maxima around 600 nm whilst the fifth spectrum features clearly one larger absorption feature between the local maxima around 550 nm and 750 nm. Thus, if your objectives include to compare the absorption in the red edge of different spectra, these two spectra would not be comparable due to the fact that this large feature is split into several smaller features in spectrum 1.

However, you might have the impression that some of the local maxima could be removed because they are very small and maybe afflicted with uncertainties which might legitimate it to manipulate the continuum line. Therefore, `hsdar` provides functions to remove and add "continuum points" to a continuum line which allows to adapt the continuum line which can be used to also adapt band depth or ratio transformation. Handle these functions with care to avoid continuum lines too much build by subjective methods.

If you have a large Speclib, its quite labor-intensive to manually adapt the continuum lines because you have to go through every sample in you Speclib. In the following example the procedure will be shown with one exemplary sample:

Continuum points can be deleted using the `deletetcp` function:

```

> str(deletetcp)

function (x, ispec, cpdelete)

```

with `x` is the continuum line, `ispec` is the name or index of the spectrum to be modified and `cpdelete` is a single value or vector of wavelength containing fix points to be deleted. Comparing spectrum 1 and spectrum 5 we have seen that the continuum line of spectrum 1 features additional local maxima beyond 580 nm which are, however, not physically explainable so that you might want to delete all points between 580nm and 700nm:

```

> gettcp(sh_cline, 1, subset = c(500, 700)) #see all points

$ptscon
  Wavelength Reflectance
1         566      10.93322
2         567      10.93838
3         568      10.94255
4         569      10.94506
5         570      10.94632
6         601      11.09275
7         604      11.10690
8         605      11.11117
9         606      11.11376
10        607      11.11505
11        636      11.29697
12        637      11.30238
13        638      11.30288

$ispec
[1] 1

> sh_cline <- deletetcp(sh_cline, 1,
+                       c(580:700)) #delete all between 580 and 700 nm
> gettcp(sh_cline, 1, subset = c(500, 700)) #see what happened

$ptscon
  Wavelength Reflectance
1         566      10.93322
2         567      10.93838
3         568      10.94255
4         569      10.94506
5         570      10.94632

$ispec
[1] 1

```

Similarly you could add a continuum point by specifying the wavelength of the point to be added. Though it doesn't make sense in this context you could add a point at the wavelength 460:

```

> #sh_cline <- addtcp(sh_cline, 1, 460)

```

After modifying the continuum line by adding and/or deleting continuum points, you can check the line for intersection with the spectrum using the checkhull function:

```

> checkhull(sh_cline, 1)$error
[1] 0 0

```

If there are any errors, additional continuum points have to be defined to meet the constraint that the hull does not cross the spectrum.

```
> sh_cline <- addcp(sh_cline, 1, 317)
```

After all uncertainties are removed, the hull can be re-calculated using "makehull":


```
> sh_clineUpdate <- makehull(sh_cline, 1) #update the hull of spectrum 1
```

After all hulls of the Speclib are modified and corrected, the transformed Speclib has to be updated with the new hulls:

```
> sh_bd <- updatecl(sh_bd, sh_clineUpdate) #update the band depth
```

Now, we can plot the resulting continuum removed spectra:

```
> #plot new line:
> par (mfrow = c(1,2))
> plot(sh_cline, ispec = 1, main = "Updated Segmented hull",
+      subset = c(500,800))
> #plot new band depth
> plot(sh_bd, ispec = 1, main="Updated hull - Band depth",
+      subset = c(500,800))
```



```
> data(spectral_data)
> spectral_data_preproc <- smoothSpeclib(spectral_data,
```

```

+                                     method = "sgolay", n = 5)
> mask(spectral_data_preproc) <- c(1040,1060,1300,1450)
> sh_cline <- transform_speclib(spectral_data_preproc,
+                               method = "sh", out = "raw")
> par (mfrow = c(2,1))
> plot(sh_cline, 1, subset = c(550, 650))
> plot(sh_cline, 5, subset = c(550, 650))
> getcp(sh_cline, 1, subset = c(500, 700))
> getcp(sh_cline, 5, subset = c(500, 700))
> sh_cline <- deletetcp(sh_cline, 1, c(550:700)) #delete
> checkhull(sh_cline, 1)$error
> sh_cline <- addcp(sh_cline, 1, 560)
> sh_clineUpdate <- makehull(sh_cline, 1)
> sh_bd <- transform_speclib(spectral_data_preproc,
+                             method = "sh", out = "bd")
> sh_bd <- updatecl(sh_bd, sh_clineUpdate)
> par (mfrow = c(1,2))
> plot(sh_cline, 1, main = "Updated Segmented hull",
+       subset = c(500,800))
> #plot new band depth
> plot(sh_bd, 1, main="Updated hull - Band depth",
+       subset = c(500,800))
>

```


6.2 Extracting absorption features

Let's continue with the preprocessed spectra from [section 6](#)

```

> sh_bd <- transform_speclib(spectral_data_preproc,
+                             method = "sh", out = "bd")
> ## Define features automatically
> features <- define.features(sh_bd)
> ##Example to isolate the features around 450,700,1200 and 1500nm.
> featureSelection <- specfeat(features, c(450,700,1200,1500))
> ## Plot features
> plot(featureSelection, 1:4)


```



Some features are larger than others. For some research questions it might be important to cut the features at specific wavelengths.

The following example shows how to cut the first two features:

```
> featuresCut <- cut_specfeat(featureSelection, fnumber = c(1,2),
+                             limits = c(c(310, 560), c(589, 800)))
> ## Plot result
> plot(featuresCut, 1:2)
```



7 Basic data analysis tools

7.1 T-Test

You can compare the distribution of two subsets of a Speclib by a T-test. At each wavelength subset one is compared to subset two: First, split "spectral_data" into a Speclib from Kailash site and a Speclib from Namco site:


```
> #split into subsets:
> sp_namco <- subset(spectral_data, Site == "Namco")
> sp_kailash <- subset(spectral_data, Site == "Kailash")
```

Then perform a T-test: The table shows the statistical parameters of the T-test (e.g. p-value, confidence intervals, mean values) for each wavelength. E.g. for the wavelength 750 nm the differences are highly significant with higher reflectance at the Kailash site (31.11%) than on the Namco site (24.99%). Let's visualize how the reflexion at the red edge reflexion shoulder (around 750 nm) of the Kailash samples compares to the reflexion at the same position of the Namco sites:


```

> boxplot(spectra(sp_namco)[,wavelength(sp_namco) == 750],
+         spectra(sp_kailash)[,wavelength(sp_kailash) == 750],
+         names=c("Namco","Kailash"), ylab = "Reflectance")

```



7.2 Regressions

Currently, `hsdar` provides no special tools for regressions, however you can use the standard R routines. E.g. if you want to test each wavelength for its relation to plant fraction, write a loop over each wavelength and access the spectra and the attribute "PV" from the `SpecLib`:

```

> result <- list()
> for (i in 1:length(wavelength(spectral_data)))
+ {
+   result[[i]] <- summary(lm(spectra(spectral_data)[,i] ~
+                             attribute(spectral_data)$PV))
+ }
> names(result) <- wavelength(spectral_data)

```

In this way you can access the results using the wavelengths as string. E.g. let's have a look on the results at 650 nm:

```

> result$"650"

Call:
lm(formula = spectra(spectral_data)[, i] ~ attribute(spectral_data)$PV)

Residuals:
      Min       1Q   Median       3Q      Max
-14.5195  -1.8714   0.3165   2.4116   7.5872

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)    26.45560    0.84730   31.22  <2e-16 ***
attribute(spectral_data)$PV -0.26777    0.01689  -15.86  <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 3.755 on 79 degrees of freedom
Multiple R-squared:  0.7609,    Adjusted R-squared:  0.7579
F-statistic: 251.5 on 1 and 79 DF,  p-value: < 2.2e-16


```

You can easily plot this relation by a scatterplot with PV on the x-axis and the reflectance at the wavelength 650 nm on the y-axis.

```

> plot(attribute(spectral_data)$PV,
+       spectra(spectral_data)[, wavelength(spectral_data)==650],
+       xlab = "PV", ylab = "Reflectance at 650nm")

```



8 Calculating spectral indices

There are three different kinds of spectral indices implemented in `hsdar`: A variety of common as well as recently developed vegetation indices, red edge parameters and normalised ratio indices.

8.1 Vegetation indices

To see the whole set of implemented indices and how they are calculated please read the `hsdar` help or manual. The indices are calculated for each sample in the SpecLib. For example calculate the NDVI like this:

```
> data(spectral_data)
> ndvi <- vegindex(spectral_data, "NDVI")
> ndvi #see ndvi
```

[1]	0.45196160	0.42158241	0.29653444	0.43702210	0.52713542	0.50245462
[7]	0.38192257	0.56776629	0.45897541	0.61313405	0.68954323	0.61009374
[13]	0.57343401	0.46233470	0.57857696	0.31392631	0.24206030	0.68671665
[19]	0.61468022	0.60937227	0.59271988	0.54379647	0.51993521	0.62576449
[25]	0.60474869	0.20507617	0.72450663	0.70809014	0.79646640	0.54337931
[31]	0.70351218	0.07998129	0.07020182	0.11105976	0.41060959	0.10986429

```
[37] 0.43882499 0.07695109 0.09516512 0.22172498 0.22641970 0.14415680
[43] 0.50426263 0.03947298 0.26349167 0.42471459 0.06067836 0.30657992
[49] 0.09462015 0.19965990 0.45751440 0.39683012 0.42437552 0.40736745
[55] 0.37949699 0.32725858 0.16914661 0.32534664 0.13632719 0.31219562
[61] 0.24381867 0.14626527 0.30131823 0.43613739 0.04693296 0.25804194
[67] 0.23836167 0.20243122 0.30541179 0.27896223 0.34706871 0.07461242
[73] 0.22197914 0.42038252 0.34904643 0.51969969 0.66462644 0.67136530
[79] 0.40905723 0.47450543 0.48437841
```

You can also directly calculate all available indices by creating a vector of the names of all already implemented indices with "vegindex()" which is used as index parameter in the vegindex function:

```
> avl <- vegindex()
> vi <- vegindex(spectral_data, index = avl)
```

8.2 Red edge parameters

Shape and location of the red edge are commonly described by four parameters:

R0 the minimum reflectance in the red spectrum

l0 wavelength of the minimum reflectance

lp inflection point


Rs shoulder wavelength

The red edge parameters are calculated as proposed in [Bach \(1995\)](#) from the spectral area between 600 and 900 nm. l0 is calculated as the last root before the maximum value of the 2nd derivation. The minimum reflectance is the reflectance at (l0). The inflection point is the root of the 2nd derivative function between the maximum value and the minimum value. The shoulder wavelength is the first root beyond the minimum value of the 2nd derivation.

```
> data(spectral_data)
> rd <- rededge(spectral_data)
```

Results can be presented as boxplot. For example create a boxplot for R0:

```
> boxplot(rd$R0 ~ spectral_data$attributes$Site, ylab = "R0")
```



8.3 Normalised ratio indices

`hsdar` has implemented a method to calculate NDVI-like Normalised ratio indices (NRI) (also named as narrow band indices). Thus for all possible band combinations in the spectrum, the following calculation is performed:

$$nri_{B1,B2} = \frac{R_{B1} - R_{B2}}{R_{B1} + R_{B2}} \quad (1)$$

with R being reflectance values at wavelength $B1$ and $B2$, respectively.

With this function you could now calculate the NRI for all band combinations, however this requires some time so that we will explain the NRI using resampled bands. Resample "spectral_data" to the resolution of WorldView-2-8:

```
> spec_WV <- spectral.resampling(spectral_data, "WorldView2-8",
+                               response_function = FALSE)
```

Now, see how nrIs are calculated:

```
> str(nri)
```

```
function (x, b1, b2, recursive = FALSE, bywavelength = TRUE)

> help(nri)
```

There are two possibilities what to do with nri. Either you could assign two bands by wavelength from which the NRI should be calculated or you can assign "recursive=TRUE" which means that NRI are calculated for all possible band combinations. For our case with 8 WorldView channels this would mean that $8 \times 7 = 56$ combinations will be calculated for each spectrum in the SpecLib.

```
> nri_WV <- nri(spec_WV, recursive = TRUE)
> nri_WV

Data: nri, dimension: 8, 8, 81

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1159411      NA      NA      NA      NA      NA
[3,] 0.2803113 0.1698916      NA      NA      NA      NA
[4,] 0.3419621 0.2353521 0.06818695      NA      NA      NA
[5,] 0.3789481 0.2750933 0.11035954 0.04249234      NA      NA
[6,] 0.6008295 0.5211952 0.38543239 0.32580814 0.2872932      NA
[7,] 0.6962423 0.6312583 0.51679015 0.46498865 0.4310124 0.1640306
[8,] 0.7144031 0.6525083 0.54278779 0.49284142 0.4599820 0.1989846

      [,7] [,8]
[1,]      NA      NA
[2,]      NA      NA
[3,]      NA      NA
[4,]      NA      NA
[5,]      NA      NA
[6,]      NA      NA
[7,]      NA      NA
[8,] 0.0361334      NA

... (79 layers omitted)

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1576740      NA      NA      NA      NA      NA
[3,] 0.3199467 0.1708939      NA      NA      NA      NA
[4,] 0.3513848 0.2050727 0.03542016      NA      NA      NA
[5,] 0.3637979 0.2186670 0.04962758 0.01423243      NA      NA
[6,] 0.5994680 0.4879117 0.34585556 0.31428549 0.3014013      NA
[7,] 0.6942123 0.6024859 0.48112968 0.45343688 0.4420573 0.1622773
[8,] 0.7038670 0.6143777 0.49550894 0.46830800 0.4571224 0.1806043

      [,7] [,8]
[1,]      NA      NA
[2,]      NA      NA
[3,]      NA      NA
[4,]      NA      NA
[5,]      NA      NA
```

```
[6,]      NA      NA
[7,]      NA      NA
[8,] 0.01888041      NA
      wavelength of length = 8
      fwhm for each wavelength
```

To get access to the NRI of a specific spectrum use

```
> str(nri_WV)
```

to see that the NRI can be accessed via `object$nri`. This object has three dimensions: band 1, band 2 and the spectrum. Therefore type the following to see the NRI for all band combinations of the first spectrum in the SpecLib:

```
> nri_WV$nri[, , 1]

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1159411      NA      NA      NA      NA      NA
[3,] 0.2803113 0.1698916      NA      NA      NA      NA
[4,] 0.3419621 0.2353521 0.06818695      NA      NA      NA
[5,] 0.3789481 0.2750933 0.11035954 0.04249234      NA      NA
[6,] 0.6008295 0.5211952 0.38543239 0.32580814 0.2872932      NA
[7,] 0.6962423 0.6312583 0.51679015 0.46498865 0.4310124 0.1640306
[8,] 0.7144031 0.6525083 0.54278779 0.49284142 0.4599820 0.1989846

      [,7] [,8]
[1,]      NA      NA
[2,]      NA      NA
[3,]      NA      NA
[4,]      NA      NA
[5,]      NA      NA
[6,]      NA      NA
[7,]      NA      NA
[8,] 0.0361334      NA
```

Note that the resulting matrix only contains the indices for one side of the matrix because the information content would be the same for the other side of the matrix only with opposite algebraic sign.

8.4 Comparing distributions of NRI

See section 9 for how to relate NRIs to environmental variables which is most likely what you want to achieve with your NRIs.

9 Analysing relations between NRI and environmental variables

This section will show how to relate the NRI (see section 8.3) to environmental variables and how to create nice plots with a lot of information content.

If you haven't already calculated the NRI from WorldView-resampled bands, do it now to work through this section:

```
> spec_WV <- spectral.resampling(spectral_data, "WorldView2-8",  
+                               response_function = FALSE)  
> nri_WV <- nri(spec_WV, recursive = TRUE)
```

9.1 Correlations

In this example we want to correlate each NRI to the fraction of photosynthetically active vegetation (PV) to determine to which degree the NRI are related to PV. Use the Speclib and NRI data created in section 9. First create a new variable from the attributes of the Speclib containing the fraction of green vegetation per sample:

```
> veg <- attribute(spec_WV)$PV
```

Then you can correlate this to the NRI:

```
> cortestnri <- cor.test(nri_WV, veg)
```

See how the output of such a correlation is printed:

```
> cortestnri  
  
Data: nri, dimension: 8, 8, 81  
      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]  
[1,]      NA      NA      NA      NA      NA      NA  
[2,] 0.1159411      NA      NA      NA      NA      NA  
[3,] 0.2803113 0.1698916      NA      NA      NA      NA  
[4,] 0.3419621 0.2353521 0.06818695      NA      NA      NA  
[5,] 0.3789481 0.2750933 0.11035954 0.04249234      NA      NA  
[6,] 0.6008295 0.5211952 0.38543239 0.32580814 0.2872932      NA  
[7,] 0.6962423 0.6312583 0.51679015 0.46498865 0.4310124 0.1640306  
[8,] 0.7144031 0.6525083 0.54278779 0.49284142 0.4599820 0.1989846  
      [,7] [,8]  
[1,]      NA      NA  
[2,]      NA      NA  
[3,]      NA      NA
```



```

[4,]      NA  NA
[5,]      NA  NA
[6,]      NA  NA
[7,]      NA  NA
[8,] 0.0361334  NA

... (79 layers omitted)

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1576740      NA      NA      NA      NA      NA
[3,] 0.3199467 0.1708939      NA      NA      NA      NA
[4,] 0.3513848 0.2050727 0.03542016      NA      NA      NA
[5,] 0.3637979 0.2186670 0.04962758 0.01423243      NA      NA
[6,] 0.5994680 0.4879117 0.34585556 0.31428549 0.3014013      NA
[7,] 0.6942123 0.6024859 0.48112968 0.45343688 0.4420573 0.1622773
[8,] 0.7038670 0.6143777 0.49550894 0.46830800 0.4571224 0.1806043


      [,7] [,8]
[1,]      NA  NA
[2,]      NA  NA
[3,]      NA  NA
[4,]      NA  NA
[5,]      NA  NA
[6,]      NA  NA
[7,]      NA  NA
[8,] 0.01888041  NA
      wavelength of length = 8
      fwhm for each wavelength
Call: cor.test(NA NA NA)

Models contain following parameters:
[[1]] p.value
[[2]] estimate
Dimension of each parameter: 8, 8, 1

```


As you can see, there are p values and estimates of the correlation stored for each band combination. The coefficients of the correlation can be visualized by the function by plotting the object. The p-value is in most cases the interesting coefficient:

```
> plot(cortestnri, coefficient = "p.value")
```



Now it becomes obvious that the NRI from the band combination 4 and 3 is the definitely not correlated to vegetation cover. To see which NRI are significantly correlated let's only plot the p-values where NRI was correlated with a p value less than 0.01:

```
> plot(cortestnri, coefficient = "p.value", range = c(0,0.01))
```



Obviously all other NRI except the mentioned NRI from the bands 4 and 3 are significantly correlated to vegetation cover.

9.2 Linear models

Linear regressions between NRI and environmental variables can be performed using the "lm.nri" function. See how this function works:

```
> str(lm.nri)

function (formula, preddata = NULL, ...)
```

The function requires the formula of the model as well as "preddata" which is a Speclib or a data.frame containing the environmental variables. Use the Speclib and NRI data created in section 9 to perform a linear regression between NRI and fraction of vegetation which is stored as attribute in "spec_WV":

```
> lmnri <- lm.nri(nri_WV ~ PV, preddata = spec_WV)
```

See how the lmnri object looks like:

```

> lmnri

Data: nri, dimension: 8, 8, 81

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1159411      NA      NA      NA      NA      NA
[3,] 0.2803113 0.1698916      NA      NA      NA      NA
[4,] 0.3419621 0.2353521 0.06818695      NA      NA      NA
[5,] 0.3789481 0.2750933 0.11035954 0.04249234      NA      NA
[6,] 0.6008295 0.5211952 0.38543239 0.32580814 0.2872932      NA
[7,] 0.6962423 0.6312583 0.51679015 0.46498865 0.4310124 0.1640306
[8,] 0.7144031 0.6525083 0.54278779 0.49284142 0.4599820 0.1989846

      [,7] [,8]
[1,]      NA      NA
[2,]      NA      NA
[3,]      NA      NA
[4,]      NA      NA
[5,]      NA      NA
[6,]      NA      NA
[7,]      NA      NA
[8,] 0.0361334      NA

... (79 layers omitted)

      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,]      NA      NA      NA      NA      NA      NA
[2,] 0.1576740      NA      NA      NA      NA      NA
[3,] 0.3199467 0.1708939      NA      NA      NA      NA
[4,] 0.3513848 0.2050727 0.03542016      NA      NA      NA
[5,] 0.3637979 0.2186670 0.04962758 0.01423243      NA      NA
[6,] 0.5994680 0.4879117 0.34585556 0.31428549 0.3014013      NA
[7,] 0.6942123 0.6024859 0.48112968 0.45343688 0.4420573 0.1622773
[8,] 0.7038670 0.6143777 0.49550894 0.46830800 0.4571224 0.1806043

      [,7] [,8]
[1,]      NA      NA
[2,]      NA      NA
[3,]      NA      NA
[4,]      NA      NA
[5,]      NA      NA
[6,]      NA      NA
[7,]      NA      NA
[8,] 0.01888041      NA

      wavelength of length = 8
      fwhm for each wavelength
Call: lm(nri_WV ~ PV)

Models contain following parameters:
[[1]] estimate
[[2]] std.error
[[3]] t.value
[[4]] p.value
[[5]] r.squared

```

```
Dimension of each parameter: 8, 8, 2
```

Each model contains the parameters known from common linear regressions. However, the dimensions of the parameters make clear that there are 8*8 models stored in the object for which the parameters are available. Imagine you have even more bands than 8 you will most likely want to find out which is the best performing model. You can do this using the function "nri_best_performance":

```
> str(nri_best_performance)

function (nri, n = 1, coefficient = "p.value", predictor = 2,
  abs = FALSE, findMax = FALSE, ...)
```

The function takes the NRI data and the linear (or generalized linear) model as input. Further "n" can be specified which is the number of best models which should be returned. The other parameters are not interesting at the moment.

In this example we want to get only the best model (n=1):

```
> nribest <- nri_best_performance(lmnri, n = 1)
> nribest

$Indices
  Band_1 Band_2
1    831    659

$Models

Call:
lm(formula = formula, data = glm_data)

Coefficients:
(Intercept)          PV
  0.067404      0.006621
```

Maybe it is interesting to see the NRI values of the best performing NRI. Use "getNRI" and the NRI data as well as the best performing NRI as input. For each sample of the Speclib, the NRI value of the best model is then shown:

```
> getNRI(nri_WV, nribest)

[1] 0.43101242 0.41149333 0.30831596 0.41752642 0.49128392 0.47063613
[7] 0.36826810 0.52796793 0.42846951 0.56473917 0.62450540 0.56048661
[13] 0.52440929 0.42556292 0.53634006 0.29608095 0.22836956 0.61582279
[19] 0.56364328 0.55574709 0.54529655 0.51162112 0.48490313 0.57428890
[25] 0.54988360 0.20292880 0.65341198 0.62821382 0.71988189 0.49429727
```

```
[31] 0.63419622 0.08936416 0.07351639 0.11540955 0.38363191 0.10947574
[37] 0.40517664 0.08020265 0.09789021 0.21373659 0.21690103 0.14341106
[43] 0.47481939 0.04361827 0.24530655 0.39367744 0.06715589 0.28227583
[49] 0.09690459 0.19062215 0.41966736 0.37261415 0.38650867 0.37191632
[55] 0.35838550 0.30218011 0.16218551 0.30181104 0.13420367 0.29369918
[61] 0.22698463 0.14889970 0.28737727 0.39703828 0.05436204 0.23999128
[67] 0.22505571 0.19773297 0.29223037 0.26683882 0.33937034 0.08544740
[73] 0.24616086 0.39744300 0.34152636 0.48558092 0.59647107 0.62268871
[79] 0.37951493 0.43188670 0.44205725
```

9.3 Generalized linear models


Calculation and plotting of generalized linear models work in the same way as the calculation of linear models in section 9.2. Note that the coefficients change because of the different models (e.g. r.squared is not available using glms).

9.4 Plot NRI models

Linear (or generalized linear) models of NRI and environmental variables can be plotted like shown in e.g. [Meyer et al. \(2013\)](#); [Mutanga and Skidmore \(2004b\)](#). Note: The plots in the cited studies based on NRI with narrow bands which were not resampled to e.g. WorldView channels like shown in this tutorial. You can easily create models and plots with narrow bands by omit the resampling in the beginning of this section 9 and using "spectral_data" instead of the resampled data "spec_WV".


The `plot.lmnri` function takes a model from NRI and predictor variables (see section 9.2) and the coefficient to be plotted as input. Start with plotting the r.squared values of the linear model from section 9.2:

```
> plot(lmnri, coefficient = "r.squared", main = "R squared")
```



For each band combination, the r.squared value of the model of NRI and the environmental variable is represented by colour. Maybe you want to limit your plot to only these band combinations whose NRI were significantly related to the environmental variable. You can do this using the "constraint" parameter. Assign your constraint as string. For example: constraint = "p.value<0.01" means that only r.squared values of models with p value less than 0.01 will be drawn:

```
> plot(lmnri, coefficient = "r.squared", main = "R squared",
+       constraint = "p.value<0.01")
> ## additionally mark the best performing model:
> mark_nri_best_performance(nrIBest, lmnri)
```



Note that the reason why the rectangle seems to overlap with the neighbouring index is that channels 7 and 8 overlap in the Worldview-2-sensor.

10 Linear spectral unmixing

Linear spectral unmixing is a method to derive the cover fractions of different materials within the footprint of multi- or hyperspectral pixels/measurements. For a detailed overview of linear spectral unmixing see e.g. [Sohn and McCoy \(1997\)](#). The algorithm in the `hsdar` package uses the code originally developed for Grass GIS by Markus Neteler. The basic concept behind linear spectral unmixing is that you provide a set of spectra (taken by a field spectrometer or multi/hyperspectral satellite sensor) and a set of spectra providing the information about the spectral properties of the pure materials which are mixed in the first set. The latter set is usually called "endmembers". Here, we will use the field spectra from Tibet and the endmember "PV" (photosynthetically active vegetation), "NPV" (photosynthetically not active vegetation), and "soil". All of the endmember spectra were taken in the field. Note that if endmember spectra taken in the field should be used to unmix satellite albedo values, a proper atmospheric correction is urgently required. Now back to the analysis: We will first load the datasets provided in the `hsdar` package and subset the datasets to the Kailash-site. Additionally, we must set the mask according to the mask in the endmember spectra:

```
> data(spectral_data)
> data(endmember_spectra)
```



```

> ## Subset both speclibs
> spec_ka <- subset(spectral_data, Site == "Kailash")
> em_ka <- subset(endmember_spectra, Site == "Kailash")
> ## Look at the mask in the endmember_spectra
> mask(endmember_spectra)

NULL

> mask(spec_ka) <- c(1045, 1055, 1350, 1706)
> ## Divide spectra of both speclibs by 100 to achieve the range
> ## [0, ..., 100]
> spectra(spec_ka) <- spectra(spec_ka)/100
> spectra(em_ka) <- spectra(em_ka)/100

```

Now, we start the unmixing approach:

```

> unmix_res <- unmix(spec_ka, em_ka)
> ## Let's have a look at the output:
> str(unmix_res)

List of 2
 $ fractions: num [1:3, 1:43] 0.109 0.0692 0.8294 0.1612 0 ...
 ..- attr(*, "dimnames")=List of 2
 .. ..$ : chr [1:3] "1" "2" "3"
 .. ..$ : chr [1:43] "1" "2" "3" "4" ...
 $ error      : num [1:43] 0.01818 0.00711 0.02488 0.00686 0.08275 ...

```

The return value of "unmix" is a list with two elements:

1. "fractions": A matrix with the unmixed fractions of each endmember in each spectrum. The different spectra are the columns and the endmembers the rows of the matrix.
2. "error": Mathematically speaking, an over-determined linear equation system is solved during linear spectral unmixing. Thus, one is only able to minimize the error during solving the system. This amount of error is returned as the euclidean norm of the error vector after least square error minimisation. Large error values may indicate that endmember spectra do not fit well to the mixed material spectra.

We can now test the performance of the unmixing approach. Therefore, we compare the unmixed fractions with the fractions estimated in the field:

```


> vd_npv <- data.frame(Predicted = unmix_res$fractions[1,],
+                      Observed = spec_ka$attributes$NPV/100)
> vd_pv <- data.frame(Predicted = unmix_res$fractions[2,],
+                      Observed = spec_ka$attributes$PV/100)
> vd_soil <- data.frame(Predicted = unmix_res$fractions[3,],
+                       Observed = spec_ka$attributes$Bare.soil/100)

```

```

> par(mfrow = c(2,2), mar = c(5,4,4,0.1))
> plot(Predicted ~ Observed, data = vd_pv, ylab = "Predicted PV",
+      xlab = "Observed PV", xlim = c(0,1), ylim = c(0,1))
> plot(Predicted ~ Observed, data = vd_npv, ylab = "Predicted NPV",
+      xlab = "Observed NPV", xlim = c(0,1), ylim = c(0,1))
> plot(Predicted ~ Observed, data = vd_soil, ylab = "Predicted Soil",
+      xlab = "Observed Soil", xlim = c(0,1), ylim = c(0,1))

```



References

- Bach, H., 1995. Die Bestimmung hydrologischer und landwirtschaftlicher Oberflächenparameter aus hyperspektralen Fernerkundungsdaten. Münchner Geographische Abhandlungen Reihe B, Band B21.
- Clark, R. N., King, T. V. V., Gorelick, N. S., 1987. Automatic continuum analysis of reflectance spectra. In: Proceedings of the Third Airborne Imaging Spectrometer Data Analysis Workshop. pp. 138–142.
- Jacquemoud, S., Baret, F., 1990. Prospect - a model of leaf optical-properties spectra. Remote Sensing of Environment 34 (2), 75–91.
- Jacquemoud, S. A., Verhoef, W., Baret, F., Bacour, C., Zarco-Tejada, P. J., Asner, G. P., Francois, C., Ustin, S. L., 2009. PROSPECT + SAIL models: A review of use for vegetation characterization. Remote Sensing of Environment 113, Supplement 1 (0), 56 – 66.
- Lehnert, L. W., Meyer, H., Meyer, N., Reudenbach, C., Bendix, J., 2013. Assessing pasture quality and degradation status using hyperspectral imaging: a case study from western Tibet. Vol. 8887. pp. 88870I–88870I–13.
URL <http://dx.doi.org/10.1117/12.2028348>
- Lehnert, L. W., Meyer, H., Meyer, N., Reudenbach, C., Bendix, J., 2014. A hyperspectral indicator system for rangeland degradation on the Tibetan Plateau: A case study towards spaceborne monitoring. Ecological Indicators 39 (0), 54 – 64.
- Lehnert, L. W., Meyer, H., Wang, Y., Miehe, G., Thies, B., Reudenbach, C., Bendix, J., Jul. 2015. Retrieval of grassland plant coverage on the Tibetan Plateau based on a multi-scale, multi-sensor and multi-method approach. Remote Sensing of Environment 164, 197–207.
- Meyer, H., Lehnert, L. W., Wang, Y., Reudenbach, C., Bendix, J., 2013. Measuring pasture degradation on the Qinghai-Tibet Plateau using hyperspectral dissimilarities and indices. Vol. 8893. pp. 88931F–88931F–13.
URL <http://dx.doi.org/10.1117/12.2028762>
- Mutanga, O., Skidmore, A., 2004a. Hyperspectral band depth analysis for a better estimation of grass biomass (*Cenchrus ciliaris*) measured under controlled laboratory conditions. International Journal of applied Earth Observation and Geoinformation 5 (2), 87–96.
- Mutanga, O., Skidmore, A. K., 2004b. Narrow band vegetation indices overcome the saturation problem in biomass estimation. International Journal of Remote Sensing 25 (19), 3999–4014.
- Sohn, Y. S., McCoy, R. M., 1997. Mapping desert shrub rangeland using spectral unmixing and modeling spectral mixtures with TM data. Photogrammetric Engineering and Remote Sensing 63 (6), 707–716.