# I ntroduction to Matlab <br> $+$ <br> Mathematical aspects of bilinear factor models (PCA and PLS) 

Frans van den Berg
fb@life.ku.dk
(last modifications September 22, 2007)
University of Copenhagen, Faculty of Life Sciences
Department of Food Science
Quality and Technology, Spectroscopy and Chemometrics Group Rolighedsvej 30 (Room T 447)
DK-1958 Frederiksberg-C
Denmark

This document + programming code is available form
www.models.life.ku.dk

## Contents

Introduction to basic Matlab with computer code

- Introduction
page - 02
- Matlab environment -02
- Matlab commands -04
- Scripts, functions and plotting - 11
- Manipulating I mages - 19

| - I mporting data from The Unscrambler, LatentiX |  |
| :--- | :--- |
| and MS-Excel into Matlab |  |

- The "path" and third party products - 23
- More reading material -27

Transparencies (+ some telegram-style explanation)

- Principal Component Analysis
- 30
- Partial Least Squares regression - 58

Practicing Matlab with factor models

- Linear Algebra / PCA -77
- Regression / PLS - 79
- NIR with temperature effects -80
- Basic Statistics - 85
- Linear Regression - 88

Transparencies Linear Algebra -91
Transparencies Basic Statistics -98

## I ntroduction

Two titles = two aims: 1) to get a quick introduction to the computer program Matlab (see www.mathworks.com); 2) to get some insight into the bilinear factor models Principal Component Analysis (PCA) and Partial Least Squares (PLS) regression, focusing on the mathematics and numerical aspects rather than how's and why's of data analysis practice. For the latter part it is assumed (but not absolutely necessary) that the reader is already familiar with these methods. It also assumes you have had some preliminary experience with linear/matrix algebra.

This introduction is based on Matlab releases $7 / 14$ (but the difference with older releases is insignificant for the scope of this introductory; exercises will work with any of the resent Matlab releases). The computer exercise material can be found at www.models.life.ku.dk.

## Matlab environment

Let us start up the Matlab environment. This will open the Matlab command editor, with the " $\gg$ " command prompt. Behind this prompt you can type in a huge number of commands that will perform a multitude of different
 (mainly numerical and graphical) operations, some of which will be explained below. This same prompt will be used in this document to mark a Matlab statement. Matlab commands are further "coded" by the italic font.
Thus, to get help, go to the Matlab windows and type
>> help
If you end this line by pressing [enter], Matlab will respond with some lines of information popping up in the Matlab window (never mind what-is-what for now).

The figure below shows the Matlab program workspace or environment (a similar "control panel" should have appeared on your computer screen already, with slight differences dependent on the version/release you are working in).


What do we have?
a) This is the window where Matlab commands are executed. You can type many different commands (e.g. "help") and Matlab will react with a response by showing the result, opening a new figure window, showing an error message if the command was incorrect, etc.
b) In Matlab information is stored in so-called data arrays. In the example figure a $3 \times 3$-matrix called " X " with random numbers between 0 and 1 is generated (command "rand" in the window).
c) The Workspace window shows all the data you have stored in memory (in this case only "X").
d) The Command History window shows all the previously entered commands. Here: first we cleared the command window ("clc" = clear console), then we created the " $X$ "-matrix.
e) When you open the program you jump to the directory "work" on the hard disk. This is where the files and data will be stored. When you start working on different projects it is a good idea to organize these as subdirectories in this "work"-root.
f) You can change subdirectories (projects!) in the Current Directory window.
g) Besides the options already mentioned there are many more possibilities which you can find in the (default computer) menu structure on top of the program window.

So working with Matlab is really simple: you just type in the right commands! Now all you have to figure out is which of the hundreds of commands you can use when, how and in what order. And that is the subject of the rest of this document.

## Matlab commands

Matlab stands for Matrix Laboratory. Hence, all the data handling is matrix-oriented. To enter the simple data table X $\gg X=\left[\begin{array}{llll}1 & 2 & 3 & 4 ; \\ 5 & 6 & 7 & 8\end{array}\right]$

You now have a matrix $X$ with two rows and four columns ( $2 \times 4$ ) in the workspace (";" here means new line/row). If you enter >> Y = [1 5; 2 6; 3 7; 4 8]
you will have a second data table $Y(4 \times 2)$ - four "new lines" of two numbers each = four rows, two columns.

To see the results type
>> who
or
>> whos


The figure shows the workspace including the dimensions of the matrices (The change in layout is a results from changing operating system).
Verify the sizes of " $X$ " and " $Y$ " from the "whos" command, and make sure you understand why the two matrices get this size from the commands shown above.

Two things: a) By simply typing the commands above you generate the new variables. In the same way you can "recycle" variables to contain different data (e.g. "Y = 3" will change " $Y$ " from a matrix (4x2) into a scalar value 3 - or special matrix of size (1x1). Notice that we did this already with X , just look at the two figures above. b) Matlab is case sensitive, meaning " $X$ " and " $x$ " are not the same variable names.

Here are three ways to get information for (almost) any Matlabcommand (here we use "who" as an example):

```
>> help who (short description of the command in the work
    environment)
>> helpwin who (description in a separate window)
>> doc who (longer description in a separate window)
```

The figure shows the window using the "doc"-command on "who".

a) The function you wanted explained is shown together with a brief description.
b) Then a series of examples shows its possible use.
c) A more official description of the do-and-don'ts for a function follows (often more text than you hope for!).
d) Like for most software an extensive search mechanism is available. This can be very helpful if you know what you want to do with your data, but you cannot remember what the right command was.

Back to the command prompt; we could have created the matrix $Y$ (size $4 \times 2$ ) from before a lot quicker
$\gg Y=X^{\prime}$;
where the '-addition means transpose-operation.
The ";" at the end of the command-line means: don't show the output on the screen. But if you want to see the new $Y$ anyway, just type $\rightarrow$ Y

If at some later stage Y becomes really big you can terminate the output to the screen with [ctrl]-[c] or [crtl]-[Break]. (If you are not convinced of its importance just type "rand(5000,100)" on the command prompt, and see what happens)

Lets do some numeric operations with these two data tables $\gg \mathrm{X}-\mathrm{Y} \quad$ (subtraction \& your very first error message in Matlab!)
$\gg X^{\prime} \quad$ (subtraction - the right way - leaving you with all zeros)
$\gg X+Y^{\prime} \quad$ (addition)
$\gg \mathrm{X} * \mathrm{Y} \quad$ (inner product multiplication, $2 \times 2$ result)
$\gg Y$ $\quad$ (outer product multiplication, $4 \times 4$ result)
$\gg \mathrm{Y} . * \mathrm{X}^{\prime} \quad$ (the element-wise multiplication, $4 \times 2$ result)
$\gg Y+3.5$ (add 3.5 units to each element in " $Y$ ")
The examples above showed how to work with entire data matrices/tables. The next examples illustrate how to extract part of a table. To retrieve e.g. the second-row first-column element from table $X$, and store it in a new variable named "c", enter the following line $\gg c=X(2,1)$;

In detail: we have our matrix $X(2 \times 4)$, we pick out the element second row (the row-index $=$ first-index between brackets) and first column (column-index = second-index between brackets), and we store the result in "c" (which should contain a 5; check)
E.g. if $X=\left[\begin{array}{lll}1 & 2 & 3\end{array}\right.$
$678] \rightarrow X(1,1), X(2,1)$ and $X(2,4)$

Matlab handles scalars as matrices of size (1x1) (use the "whos" command to verify this). A variable named "ans" might have popped up. It contains the answer of the last question you asked Matlab without explicitly specifying where to store the result. Thus entering $\gg X(2,3)$
will show ans $=7$; it is bad practice to rely on this variable! Making ones - with a meaningful name - can save you a lot of trouble!

To extract a full (row or column) vector from a table you can use the colon ":"-command. E.g.
$\gg y 1=Y(:, 1) ; y 2=Y(:, 2)$;
to store the first and second column of table Y in (column vectors) y 1 and y 2 .

How does it work? The ":"-separator is used to automatically generate a list of numbers, e.g.
>> 1 : 4
You will see the list of numbers (1 to 4, counter 1). Also try the following command
>> 1:3:16
$\gg 1: 2: 16 \quad$ (Where did 16 go?)
>> 15:-1:1
So if we want to have the first three rows from $Y$, second column $\gg Y(1: 3,2)$

There are two special versions of this command: ":" by itself means take all entries in this direction (thus entire row or column, like we did in the example); you don't have to specify the length. The command "end" can be used to go to the end of a row or column, thus >> $\mathrm{Y}(2$ : end, 1 )
shows the rows from 2 until the last one, first column in table Y .
One last method to get elements from a matrix by defining new variables/vectors
>> row_index = [1 3]; col_index = [1 2]; (Make new variables)
>> Y(row_index,col_index)
Notices: a) you can use every variable name you like - including numbers - but smart names make Matlab-live a lot easier; b) a variable name cannot start with a number, and you should not give it the name of a build-in Matlab command such as "whos" (this will override the command, so you cannot use it any longer).

Higher dimensional data arrays in Matlab are created and addressed similar to vectors and matrices. To produce a ( $3 \times 3 \times 2$ ) cube $Z$ enter the following code

```
>> Z(:,:,1)=[1 2 3;4 5 6;7 8 9];
>> Z(:,:,2)=[11 12 13;14 15 16;17 18 19];
```

(first slab)
(second slab)

The cube " $Z$ " is symbolized in the figure.


Indexing in this 3D data structures follows the same conventions as before: first-index = row-index (x-direction), second-index = columnindex ( $y$-direction) and third-index $=$ so-called tube or tubular-index (z-direction; "into the paper"). Higher order data tables (4D, 5D, etc.)
are hard to visualize (or imagine) but generating them in Matlab is simple: just introduce a fourth, fifth, etc. dimension in the indexing

To see the center tube (numbers 5 and 15) in this cube enter $\gg Z(2,2,:) \quad$ (second row, second column, all slabs)

Another important command related to higher order arrays is "reshape". To transform our old matrix $X(2 \times 4)$ into a small cube $C$ ( $2 \times 2 \times 2$ ) enter the following >> C = reshape(X,2,2,2)

This reshaping-operation can only work if the number of entries in the source ( $\mathrm{X}: 4 \times 2=8$ ) and target ( $\mathrm{C}: 2 \times 2 \times 2=8$ ) are the same. Notice that "reshape" is a column wise operation
$\gg$ reshape $(X, 2,2,2) \quad$ (create a cube form $X$ - check the numbers)
$\gg$ reshape $\left(X^{\prime}, 2,2,2\right) \quad$ (create a cube from its transposed - check)
One more useful command for data cubes is "squeeze":
$\begin{array}{ll}\gg Z(:, 1,:) & \text { (selects the first slab "in the paper") } \\ \gg \text { squeeze( } Z(:, 1,:)) & \text { ("squeezes" this data in a regular/flat matrix) }\end{array}$

So, the basic form to store data is in arrays of so-called double precision real numbers: scalar in size (1x1), vectors in size (1xn) for rows or ( $n \times 1$ ) for columns, matrices in size ( $n \times m$ ), cubes in size (mxnxo), etc. There are many more forms for storing numerical and non-numerical data, important for computer-internal business, but that is only important for more advanced work.

Briefly three structures you could use or encounter. Keep a close eye on the "Workspace" area of your Matlab surrounding to see if you understand all the symbols appearing. To form a string of characters (so-called char array) enter
>> surname = 'van den Berg' (make a text string)
>> forename $=$ 'Frans'
(and another one)
>> name = [surname ', ' forename] (manipulate strings like arrays)
The so-called "cell array" can be used to store both text and numerical arrays with different sizes (this is a rare/advanced format, but notice that e.g. rows of different length cannot be stored in a conventional matrix; also note the curly brackets)
>> phone_nrs = [3528 3545]
$\gg$ CPR_nr $=2003680000$
>> nr = \{phone_nrs CPR_nr 'my numbers'\}

The last data format - "struct" or structure - is convenient to group like-things together (notice the use of "." - dot).
>> person.email = 'fb@kvl.dk' (a new text string)
$\gg$ person. name $=$ name $\quad$ (an existing text string)
$\gg$ person.numbers $=\mathrm{nr} \quad$ (cell array in the structure array)
>> person
(shows the contents of the structure)
$\gg$ person.email (to address elements use dot-notation)


Some other useful command

```
>> clear y1 y2 (clear variables y1 and y2 from the workspace)
>> save myfile (saves all variables in the workspace to a file called
>> clear all
>> load myfile
(clear variables y1 and y2 from the workspace)
(saves all variables in the workspace to a file called
"myfile.mat")
(clears all variables)
(loads all variables from the file "myfile.mat")
```

There is one more convenient feature in the Matlab editor: the arrow-up [ $\uparrow$ ] of the keyboard will callback all the old command entries, from newest to oldest. So you don't have to type everything over and over again, just call an older version of a similar thing you want to do by repeatedly pushing [ $\uparrow$ ], and edit it to do the new job. Use $[\leftarrow],[\rightarrow]$, [Delete], etc. on the command line like you would in a conventional text editor.

If you know the first letter of the old command, type it at the prompt and push the arrow-up to scroll the commands starting with this letter. E.g. typing "w" and then arrow-up would first show the "whos" and then "who" for the commands treated so far in this text.

## Script, functions and plotting

To fully explore the graphical possibilities of Matlab we first have to get hold of some interesting things to plot. We will generate some Gaussian-curves. Below is the formula for a gauss- or normaldistribution where " $x$ " is the running index (a probability index in statistics), " $\mu$ " is the mean and " $\sigma$ " is the standard deviation:

$$
y(x)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{1}{2}\left(\frac{(x-\mu)}{\sigma}\right)^{2}}
$$

We can implement and plot this in Matlab with the following lines of code

```
>> x = 1:100; mu = 50; sigma = 10;
                            (running index, mean and std. deviation)
>> y = 1/(sigma*sqrt(2* pi))* exp(-0.5*((x-mu)/sigma).^2);
    ("sqrt" = square root; "exp" = exponential)
```

Say you would like to plot green dot-markers for the points in our vectors $x$ and $y$ together with a blue line connecting the points (the default plot), enter the following command
$\gg \operatorname{plot}\left(x, y, x, y,{ }^{\prime} \cdot g^{\prime}\right) \quad$ (plot the results twice: blue line and green dots)
$\gg$ grid (add some lines to assist in interpretation)
>> xlabel('x-index'); ylabel('probability'); title('Gauss')
>> legend('line','markers') (add some text and labels for the lines)
Notices: a) If you want to pass text to Matlab you have to include it in quotes (e.g. 'x-index') just like we did for string arrays (this is the way the program distinguishes commands and variables names from textinput); b) If you do not specify markers, points will be connected by a line; If you do not specify the color the fixed order of lines will be: blue, green, red, cyan, etc.; c) study the help for the "plot" routine to get more info.

You noticed that the plot-command opened a separate window, copied in the figure below.

a) The data is drawn. The first time as a blue line (default since no plotting mode was specified) and second time green "."-symbols specified in the plot command.
b) Labels and title are next to the axis.
c) A legend box is put inside the plotting-paper.
d) A number of operations can be performed on the plot: select objects, add text/arrow/lines, zoom in/out or rotate the figure (we will come back to rotation later on).
e) Under the menu options [Edit] and [Tools] you will find more (albeit more difficult) figure manipulation options.

The commands you typed in to generate the gauss-curves are pretty complicated, and it is easy to make errors when typing complicated commands on the Matlab prompt. To avoid this you can use so-called scripts. Push the [new]-button on the Matlab window (blank in upper left corner). This will open a new window - the script editor.

In the editor sheet you can enter the following text (copy-paste, but remove the " $\gg$ " signs)
>> clear all
$\gg x=1: 100 ;$ \% running index
$\gg \mathrm{mu}=50+$ randn(3,1)*10; sig = $10+$ randn $(3,1) * 2$;
$\gg$ \% just some random numbers
$\gg y(1,:)=1 /(\operatorname{sig}(1) * \operatorname{sqrt}(2 * \mathrm{pi})){ }^{*} \exp \left(-0.5^{*}((x-\operatorname{mu}(1)) / \operatorname{sig}(1)) .{ }^{\wedge} 2\right)$;
$\gg y(2,:)=1 /(\operatorname{sig}(2) * \operatorname{sqrt}(2 * \mathrm{pi})) * \exp \left(-0.5^{*}((x-\mathrm{mu}(2)) / \operatorname{sig}(2)) . \wedge 2\right)$;
$\rightarrow>y(3,:)=1 /(\operatorname{sig}(3) * \operatorname{sqrt}(2 * p i)) * \exp \left(-0.5^{*}((x-m u(3)) / \operatorname{sig}(3)) .{ }^{\wedge} 2\right)$;
$\gg$ figure \%opens a new window
$\gg \operatorname{plot}(x, y)$; \% plot the results
$\gg$ grid;
>> legend('g1','g2','g3'); \%that's it!!!
The command "randn" generates a random number with a so-called normal/Gaussian distribution; see "helpwin randn" and "helpwin rand" for further information.

Then save this file under the name "myscript". Matlab will store it as the file called "myscript.m", and you can execute this series of commands by simply typing
>> myscript

a) Regular Matlab commands within the script file.
b) Comment text lines (always starting with "\%") help to explain the intension of the operations. They are neglected by the Matlab command interpreter.
c) Color schemes in the editor improve the readability.
d) Line numbers are useful for the error messages generated in the Workspace Command Window.

Note: scripts also are a great help and time-saver when the computations of yesterday looked a lot better than the "same" computations today, so use them!

Two related command:
>> type myscript
will print the contents of the file in the command window, and $\gg$ open myscript (first close the editor for the best effect!)
will open a script in the text editor.
Scripts are the perfect tools for work in progress, but for an operation that you do on a daily basis so-called "functions" are more appropriate.
To load the function "gauss.m" into the editor type (make sure the computer exercise material from www.models.life.ku.dk $\rightarrow$ Teaching $\rightarrow$ Introduction to Matlab + ... - bottom of the page, unzip file - is in the Matlab "work" directory)
>> open gauss

```
5) C:ttemplgauss.m
File Edit view Iext Debug Breakpoints Web window Help
```



```
a 1 }->\mathrm{ function Z = gauss ( }n,\mathrm{ mu,sigma,options)
    * function Z = gauss(n,mu,sigma,options)
    030407 FvdB
    Computes Gaussian profiles.
    * in: n (l x l) number of points in z
    mu (objects x 1) first moments of gaussian distributions in units 'points'
    sigma (objects \times 1) second moment of gaussian distributions in units 'points'
    options (1 x l) trigger N(0,s)-noise addition of fraction (default = 0 = no noise; e.g. 0.1 adds 10*-max-value noise)
    * out: Z (objects x n) data table with gaussian curves
    if (nargin < 3) }\leftarrow\mp@subsup{\mathbf{e}}{\mathbf{1}}{
        help gauss
        return
    elseif (nargin == 3)
        options = 0;
    end}\quad\leftarrow\mp@subsup{\mathbf{e}}{\mathbf{1}}{
    nz = length(mu);
        nz~= length(sigma) }\quad\leftarrow\mp@subsup{\mathbf{e}}{2}{
        error('Error: number of entries for first (mu) and second moment (sigma) must be the same')
    end
    x = l:n;
    for a=1:nz }\leftarrow\mp@subsup{\textrm{P}}{3}{
        z(a,:)}=\operatorname{exp}(-0.\mp@subsup{5}{}{*}((x-\mathrm{ -uи(a))/signa(a)).^2)./(signa(a)*(2*pi)^0.5);
end
if options(1) }\leftarrow\mp@subsup{\mathbf{C}}{4}{
    max = max (Z(:))*options(1);
    z=2+randn(size(z))*max;
end
\leftarrow e _ { 4 }
```

In the window you can see the typical structure of a Matlab function:
a) The function description with the function call name ("gauss"), input parameters ("n", "mu", "sigma" and "options") and output parameters ("Z"). Input is the data you pass from the Matlab workspace to the function, output is what the function returns to the workspace as answer to the function call. Note that the names proposed in the function header are only for explanation and internal use. You can pass any variable you like. So the call command $X=$ gauss(hundred,averages,spread) is perfectly legal. It calls the function gauss, translates the real to virtual variables (e.g. " $n$ " = "hundred"), and executes the commands. Note also that input and output parameters are optional. You could e.g. define no output parameter, so call gauss(hundred, averages,spread). Then the function would compute the curves and the result would go ... well nowhere (not very useful, but perfectly legal). We already made use of this for the "plot" command: nothing changed in the workspace, but we got a figure, so it can be very useful.
b) The help-text just after the function header should explain what to fill in to make the function perform. This is also the text you get on the screen when you type "help gauss".
c) The next step is usually some error tracking. In this example: check if enough input parameters were entered for the function (3), and see if the number of $\mu$ 's and $\sigma$ 's are the same (every Gaussian distribution requires one of each).
d) Then the actual code - the hart of the function - is executed: compute distributions of length $n$-points for each entry in "mu" (and thus "sigma"), and store the result in " $Z$ ".
e) Archetypal for the internal of functions is computer programming flow-controls. We will not explain these programming tools in detail in this introductory course, but the three flow-controls in the figure are very intuitive: (1) check ("if ... end") to see if the user passed the three obligatory input arguments ("nargin"), and if not show help end terminate the function ("break"); (2) check (second "if ... end") to see if "mu" and "sigma" are of equal length, and if not terminate the function showing an error message in the workspace ("error"); (3) loop ("for ... end") through the Gaussian distribution computation "nZ" times, and store the result in rows of matrix "Z"; (4) if input argument "options" is activated (fill in "1") some random noise will be added to the gauss-curve.

The important difference between scripts and functions is the way variables are handled. Every time a function call is made a new (small) workspace is created with all the variables necessary to execute all the commands inside the function. Thus, variable names inside the function and in the workspace are separated. If e.g. you have a variable " $x$ " in your Matlab window, and you call the gauss-function, which creates its own " $x$ " (line 25), the contents of " $x$ " in your own workspace will remain the same. This is not the case for script or command-line calls, where the " $x$ "-variables will be overwritten. Functions make a new work space, execute the required operation, clear all the temporary variables, and drop the output variable in the Matlab root-workspace (nice and clean!).

You will probably not write many functions yourself at this stage, but all the computations you will make (and actually many of the Matlab functions used so far, just type at the prompt open('mean')) are function-calls. So, you have to know what to give as input and output.
Another example:
>> myx = - pi:.1:pi; (index running from $-\pi$ to $\pi$ with a step

```
>> mymu = 31+myx*8;
>> mysig = cos(myx)*2+8; (some standard deviations)
>> myX = gauss(length(mymu),mymu,mysig); (make a data table)
```

Notices: a) "length" returns the length of a vector; b) try to figure out the size of "myX" with the command "size"; c) again, Matlab is case sensitive in variable names ( $m y x \neq m y X$ )!

Now we have the data to do some advanced plotting (just try and see; with the [U]-button in the figure window you can rotate plots to get a better view):

```
>> figure (opens an empty figure window)
>> plot(myx,myX)
>> mesh(myX)
>> surfc(myX); shading interp; colorbar hort
>> subplot(2,2,1) ("subplot" = small plots in figure window)
>> contour(myX)
>> subplot(2,2,2)
>> contourf(myX)
>> subplot(2,2,3)
> meshc(myX)
>> subplot(2,2,4)
>> pcolor(myX); shading interp; colorbar vert
```



Just a few more commands that might come in handy:
$\gg$ what (shows the ${ }^{*}$.mat and *.m file in the directory)
$\gg$ which name (tells which file "name" you are using)
$\gg$ dir (shows all the files in the working directory)
$\gg$ cd (to change directories a-la dos-style)
>> tic; pause(5); toc (these are really three functions: "tic" starts a stopwatch, "pause(5)" holds all operations for 5 seconds, while "toc" stops the stopwatch and shows the time elapse: a convenient way to time how long a Matlab function takes)
>> $x=r a n d(5,1) ; ~ m e a n(x) ; ~ s t d(x) ; ~ v a r(x) ; ~ s q r t(x) ; ~ \log (x) ; ~ \log 10(x) ;$ $\exp (x)$; (Matlab has of course a whole range of arrhythmic operators)
A few special elements:
$\begin{array}{ll}\gg a=[] & \\ \gg b=\text { NaN } & \\ & \begin{array}{l}\text { (creates an empty or (0xigns the "number" not-a-number to } b, \text { used to represent } \\ \\ \gg c=\text { missing values inside a Matlab matrix) }\end{array} \\ & \\ \text { (Assigns the "number" infinite }\end{array}$
$\gg \mathrm{c}=\mathrm{Inf} \quad$ (Assigns the "number" infinite to c , useful in logical operations)

## Manipulating I mages

To get a little more practice in manipulating data tables we will import an image into matlab. (Notice that an image is nothing more then a multi-way array!)

First import and show the image into a variable called " A " (check the workspace)
>> A = imread('apple_green.jpg');
>> image(A);
We just loaded an image of green apples, 165 by 213 pixels. Color images are most often coded by three values Red, Green and Blue (socalled RGB-indexing). Hence the data cube is ( $165 \times 213 \times 3$ ). The class of variables in this image is a special one called "uint8" or unsigned-integer-eight-bits which can hold values between whole 0 and 255.


```
Try the following commands
>> imagesc(A(:,:,1)) (only show the R-value/slab - notice sc = scale)
>> colormap gray (use gray-scale colormap)
>> imagesc(A(:,:,2)) (only show the G-value/slab)
>> imagesc(A(:,:,1)) (only show the B-value/slab)
>> image(A(:,:,[\begin{array}{lll}{3}&{1}&{2}\end{array}]))\quad("false-images" - rearranging R,G and B )
>> image(A(50:130,80:160,:)) (zoom in on one apple)
>> mesh(double(A(:,:,1))) (3D relief of the R-value matrix)
>> image(uint8(abs(double(A)-255))) (never mind, just try it!)
```

From a data point of view images are nothing else then other matrices or data tables/cubes!

## I mporting data from The Unscrambler, LatentiX and MS-Excel into Matlab

Many advanced ways for importing data into the Matlab exist (including low-level C-like programming). Most of these are too complicated for this introductory course, but we will try two routes importing external information from The Unscrambler and MS-Excel, since experience has shown that many course people have their data in these formats.

Make a small dummy-set that looks like this in The Unscrambler (don't forget the object and variable names):


Next, go to the menu [File] $\rightarrow$ [Export], then select a file-name (a) and "Matlab Export" as type (b):


If subsets of Samples (a) or Variables (b) have been defined in the Unscrambler analysis, you can select which groups to export:


In Matlab the dummy-set will look like this, with (a) objects labels, (b) variable labels and (c) the data:


Importing data into Matlab from the data analysis program LatentiX (www.latentix.com) is straightforward since the information is stored
in a compatible format with files extension "Ixf". So, to read the data from a file called "data.lxf" and store the matrix in a variable "X" enter the following commands
>> load -mat data.lxf
>> X = DataSet.data;
>> ObjLabels = DataSet.Labels\{1\};
$\gg$ VarLabels $=$ DataSet.Labels $\{2\}$;

There is a special Matlab function to import MS-Excel spreadsheets. The same dummy-set would look like this (make sure you only have one sheet to avoid problems - (a)):
$\boxed{Z}$ Microsoft Excel - data
X] Ele Edit yiew Insert Format Iools Data w



| A1 |  | $f_{x}$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | A | B | C | D | E |
| 1 |  | var1 | var2 | var3 |  |
| 2 | obj1 | 1 | 2 | 3 |  |
| 3 | obj2 | 4 | 5 | 6 |  |
| 4 |  |  |  |  |  |



In Matlab we can import the spreadsheet-file (call it "data.xls") via the following lines (a), where the data will be stored in $\mathbf{X}$ (b) and the none-numerical part in $\mathbf{N}$ (c):


Be aware that it might take some puzzling to figure out where the numbers are in the Matlab matrices, especially if text and number are intertwined in the spreadsheet. The easiest route is to make new spreadsheets (e.g. one with numbers, one with object labels, one with variable labels, etc.) that are trivial to read.

Notice: starting Matlab 7/R14 there is also a function to write to Excel files called "xlswrite". We will not explain it here, but it might be good to know for later use.

## The "path" and third party products

We already have encountered many Matlab commands. A way to see where the source of the command can be found is as follows:
>> which rand
the message on the screen shows that "rand" is a so-called build-in function.
>> which surf
now the screen message shows that "surf" is a function that can be found in "x:\MATLABXXX\toolbox\matlab\graph3d\surf.m", part of the 3D-graphics function-set.
Next to the internal function Matlab comes with a number of default toolboxes. They consist of hundreds of different functions, organized in conceptually meaningful groups. You can also buy additional toolboxes for special, very diverse purposes (e.g. statistical computations, or instrumental data input-output over the hardware computer interfaces). Even better: many scientists make there code available for free. This means you can do a lot of research by lending other peoples program code, without going into programming details (just search the internet!). However, you still have to understand how to communicate with other people's code via input and output parameters. In this paragraph we will treat organizing new features in the form of toolboxes added to the Matlab search path using the iPLS (intervalPLS) from the iToolbox written by Lars Nørgaard.

You can find the source code for the iToolbox at the KVL Quality\&Technology homepage under:
www.models.life.ku.dk $\rightarrow$ Research $\rightarrow$ Algorithms and code $\rightarrow$ iToolbox
Here you can find important information on the author, ownership and use. Download the zip-file to your own computer. Make a new directory in the path "Matlab\toolbox" called "itoolbox", and extract the contents of the zip-file to this directory. The next step is to add the new toolbox to the Matlab search path.
In the workspace editor select: [File] $\rightarrow$ [Set Path] $\rightarrow$ [Add Folder] $\rightarrow$ find [toolbox] and then [itoolbox] in the browser $\rightarrow$ [Save] $\rightarrow$ [Close] (see figure below for details).


To check if it works type
>> which ipls
$\gg$ help ipls
This will write a small help text for the toolbox in the workspace window. Can you figure out what the input and output to this function should be? We will show by illustration. The toolbox comes with an example data file called "nirbeer" (Near Infrared spectra of beer samples):
>> clear
$\gg$ load nirbeer
$\gg$ whos
(clear the workspace)
(load data)
>> M=ipls(Xcal,ycal,8,'mean',40,xaxis,'full',40);
(compute the model, 8 factors, mean centering, 40 intervals, full cross-validation and store the result in structure " M ")
>> iplsplot( M, 'intlabel')
(plot the results with numerical labels)
We will not go into the interpretation of this model. If you want to learn more about the iPLS method you can run the iPLSdemo-routine and study the manual included in the toolbox.

A few important notes: a) You have much liberty in selecting names for variables and functions which sometimes leads to problems; the order is roughly: variables in the workspace, function-files in the working directory and function-files in top-to-bottom order in the directories of the path; b) To keep things in perspective you have to be clear in naming functions and variables; c) If you get strange/unexpected error messages, use the "which" command to see which function Matlab is using, if you still get strange errors you can always try the "why" command.

As shown above, finding the right input to other peoples functions - of course depending on the help available and the didactic skills of the author - can be pretty complicated. One (very luxurious albeit time consuming way) to circumvent problems is to write a Graphical Users Interface (GUI). In short: GUI's control the user's interaction with a toolbox/set of functions by limiting the input to predefined values and a restricted (graphical) output.

One last helpful command is "lookfor". It will scan all the functions on the "path" for a particular text fragment. E.g. if you forgot the command for standard deviation enter the following line (beware, searching all the function files might take some time) >> lookfor deviation

## More reading material

This is all the basic stuff you need to know about Matlab. If you want more (introductory) information a pdf-file called "Getting Started with MATLAB" (getstart.pdf) can help out. The following parts are of interest for this introductory course (see list below for details):

1. What is Matlab (p.12)
2. Matrices and Arrays (p.19-27)
3. Basic Plotting Functions (p.88-111)
4. Flow Control and Other Data Structures (p.126-143, though not every detail!). assigned


In the next part on bilinear factor models we will only use a few specific, additional functions.
(a : important; b:less important; c: not important for this course)

## 1 Introduction

| What Is MATLAB? | 1-2 | a |  |  |
| :---: | :---: | :---: | :---: | :---: |
| The MATLAB System | 1-3 |  |  | c |
| MATLAB Documentation | 1-4 |  |  | c |
| MATLAB Online Help | 1-4 |  |  | c |
| Starting and Quitting MATLAB | 1-6 | a |  |  |
| Starting MATLAB | 1-6 | a |  |  |
| Quitting MATLAB | 1-6 | a |  |  |
| MATLAB Desktop | 1-7 | a |  |  |
| 2 Matrices and Arrays |  |  |  |  |
| Matrices and Magic Squares | 2-2 | a |  |  |
| Entering Matrices | 2-3 | a |  |  |
| sum, transpose, and diag | 2-4 | a |  |  |
| Subscripts | 2-6 | a |  |  |
| The Colon Operator | 2-7 | a |  |  |
| The magic Function | 2-8 | a |  |  |
| Expressions | 2-10 | a |  |  |
| Variables | 2-10 | a |  |  |
| Numbers | 2-11 | a |  |  |
| Operators | 2-11 | a |  |  |
| Functions | 2-12 | a |  |  |
| Examples of Expressions | 2-13 | a |  |  |
| Working with Matrices | 2-14 | a |  |  |
| Generating Matrices | 2-14 | a |  |  |
| The load Function | 2-15 | a |  |  |
| M-Files | 2-15 | a |  |  |
| Concatenation | 2-16 | a |  |  |
| Deleting Rows and Columns | 2-17 | a |  |  |
| More About Matrices and Arrays | 2-18 | a |  |  |
| Linear Algebra | 2-18 | a |  |  |
| Arrays | 2-21 | a |  |  |
| Multivariate Data | 2-24 | a |  |  |
| Scalar Expansion | 2-25 | a |  |  |
| Logical Subscripting | 2-25 | a |  |  |
| The find Function | 2-26 | a |  |  |
| Controlling Command Window Input and Output | 2-28 | a |  |  |
| The format Function | 2-28 | a |  |  |
| Suppressing Output | 2-30 | a |  |  |
| Entering Long Statements | 2-30 | a |  |  |
| Command Line Editing 3 Graphics | 2-30 | a |  |  |
| Overview of MATLAB Plotting | 3-2 |  | b |  |
| The Plotting Process | 3-2 |  |  | c |
| Graph Components | 3-5 |  |  | c |
| Figure Tools | 3-7 |  |  | c |
| Arranging Graphs Within a Figure | 3-13 |  |  | c |
| Selecting Plot Types | 3-14 |  |  | c |
| Editing Plots | 3-16 |  |  | c |
| Plot Editing Mode | 3-16 |  |  | c |
| Using Functions to Edit Graphs | 3-19 |  |  | c |
| Examples - Using MATLAB Plotting Tools | 3-20 |  |  | c |
| Modifying the Graph Data Source | 3-27 |  |  | c |
| Preparing Graphs for Presentation | 3-29 |  |  |  |
| Modify the Graph to Enhance the Presentation | 3-30 |  |  |  |
| Printing the Graph | 3-33 |  |  |  |
| Exporting the Graph | 3-35 |  |  |  |
| Basic Plotting Functions | 3-38 | a |  |  |
| Creating a Plot | 3-38 | a |  |  |
| Multiple Data Sets in One Graph | 3-40 | a |  |  |
| Specifying Line Styles and Colors | 3-41 | a |  |  |
| Plotting Lines and Markers | 3-41 | a |  |  |
| Imaginary and Complex Data | 3-43 |  |  | c |
| Adding Plots to an Existing Graph | 3-44 |  | b |  |
| Figure Windows | 3-46 | a |  |  |
| Multiple Plots in One Figure | 3-46 | a |  |  |
| Controlling the Axes | 3-48 | a |  |  |
| Axis Labels and Titles | 3-49 | a |  |  |
| Saving Figures | 3-51 | a |  |  |
| Mesh and Surface Plots | 3-52 | a |  |  |
| Visualizing Functions of Two Variables | 3-52 | a |  |  |


| Images | 3-58 | a |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Reading and Writing Images | 3-59 | a |  |  |
| Printing Graphics | 3-60 |  | b |  |
| Handle Graphics | 3-62 |  |  | c |
| Using the Handle | 3-62 |  |  | c |
| Graphics Objects | 3-63 |  |  | c |
| Setting Object Properties | 3-65 |  |  | c |
| Specifying the Axes or Figure | 3-69 |  |  | c |
| Finding the Handles of Existing Objects | 3-70 |  |  | c |
| Animations | 3-72 |  |  | c |
| Erase Mode Method | 3-72 |  |  | c |
| Creating Movies | 3-74 |  |  | c |
| 4 Programming |  |  |  |  |
| Flow Control | 4-2 |  | b |  |
| if, else, and elseif | 4-2 |  | b |  |
| switch and case | 4-4 |  | b |  |
| for | 4-5 |  | b |  |
| while | 4-6 |  | b |  |
| continue | 4-6 |  | b |  |
| break | 4-7 |  |  | c |
| try - catch | 4-8 |  |  | c |
| return | 4-8 |  |  | c |
| Other Data Structures | 4-9 | a |  |  |
| Multidimensional Arrays | 4-9 | a |  |  |
| Cell Arrays | 4-11 | a |  |  |
| Characters and Text | 4-13 | a |  |  |
| Structures | 4-16 | a |  |  |
| Scripts and Functions | 4-20 | a |  |  |
| Scripts | 4-21 | a |  |  |
| Functions | 4-22 | a |  |  |
| Types of Functions | 4-24 | a |  |  |
| Global Variables | 4-26 |  |  | c |
| Passing String Arguments to Functions | 4-27 |  |  | c |
| The eval Function | 4-28 |  |  | c |
| Function Handles | 4-29 |  |  | c |
| Function Functions | 4-30 |  |  | c |
| Vectorization | 4-32 |  |  | c |
| Preallocation | 4-33 |  |  | c |
| 5 Creating Graphical User I nter <br> 6 Desktop Tools and Developm |  |  |  | c |

## Principal Component Analysis

The next part of this document contains the transparencies used in the lectures on PCA, accompanied by some "telegram-style" notes on the main issues.

> Principal Component Analysis
> "Exploratory data analysis"

## Investigation and Visualization!




Often: observational data Sometimes: design data Seldom/never: hard models


We will use a small example as a reminder: socio-economical status of 49 countries world-wide determined by 7 indicator variables/demographic characteristics:

1) Infant death per 1000 live births
2) \# of inhabitants per physician
3) population per square kilometer
4) population per 1000 hectares of agricultural land
5) Percentage literate of population aged 15 years and over
6) \# of students enrolled in higher education per 100000 population
7) gross national product per capita (US dollars)

Notice that we don't go blind into the data: we have "independent measurements" in the form of numbers, but we also know about the sample (countries). The knowledge is important for interpretation.
R.F. Gunst and R.L. Mason "Regression Analysis and Its Application: A Data-Oriented Approach" Marcel Dekker, New York, p. 358(1980)


Principal Component Analysis (PCA) can be used to analyze tables of data, where a number of the same variables are measured/collected on a set of samples (called objects). The aim is to find the important things happening in the table: separate the main directions/principal components form the noise.
Stated differently: extract what objects have in common from what makes objects unique.


- Principal Component Analysis determines factors from a data table by making new 'pseudo'-variables (so-called principal components) from linear combinations of the original variables
- The factors (principal components) are selected to explain as much information (variance) in the table as possible
- The new variables eliminate redundant information (and filter noise) from the original data table

There are many different objectives for doing PCA - often different viewpoints associated with the specific fields of science - but will focus our attention on these three.


As a small example we will take a matrix $\mathbf{X}$ of eight object and three variables ( $8 \times 3$ ). We can visualize this table as eight points in a three dimensional space. The objects have a particular score on the three variables, and the position of an object is fixed by its three values in the Cartesian coordinate system.

In the same line we can picture (in our mind!) an eight dimensional space where the three variables (three points) are expressed in the objects-space (it is a lot harder to visualize this, but will use this mental picture later on!).


The first factor in PCA is:

- The direction that "spans" the objects the best
- Explains the maximum amount of variance in the data set
- Finds the common ("most common") direction in the set (Three times the same thing!).

To fixate the new direction in the original variable space we define three loading-values (traditionally called $\mathbf{p}$ ) plus the origin: two points to fix a line.


By projecting the sample-points on the new coordinate line (determined by the loadings), every object gets a score value (traditionally called $\mathbf{t}$ ). The $\mathbf{t}$ and $\mathbf{p}$ vector pair form the first factor (or principal component). The outer product of $\mathbf{t}$ and $\mathbf{p}$ has the same number of entries as the original matrix, and this "reconstruction" forms the best rank one (outer product) approximation of $\mathbf{X}$.


In the same way a second factor can be determined, where we select the new coordinate perpendicular to the first one (and again going through the origin).

We can continue in extracting factors until all the systematic variance is captured and only noise/information on the individual object level remains.




A bi-plot of the two main directions/principal components (total variation explained 74\%): size of the country and Poor versus Rich "Outliers" are objects that do not fit the general behavior due to errors, or in this case, due to abnormal behavior ("city-states"). Hong Kong and Singapore dominate the solution. The user often has to verify what is abnormal!


By removing the two outliers we get a better view of the same split-up: small but rich or rather, crowded but rich.


We can also remove the "crowdedness" direction by removing the appropriate variables from the data table. The new split: poor (health related) versus rich based on (American) education system.


Important: the interpretation is for the user. The labeling/coding of objects is very important (visualization!). E.g. code in Industrial, Unindustrialized and Eastern Europe (remember: the data is from 1980).

Another important issue: random reordering of rows or columns does not affect the results (only look at variances). As a user you have to keep track of the "secondary" structure (e.g. a time series, geographical or political ordering).

## Principal Component Analysis Diagnostics

- Score values $\mathbf{t}$

Position of objects in the new factor-space

- Loading values p

Role of original variables in determining the new factor-space

- Percentage explained variance

Closeness between factor space (reconstruction) and original data space

- Leverage $h$

How important is an object compared the rest of the data set $\mathrm{h}_{\mathrm{i}, \mathrm{a}} \approx \Sigma \mathrm{t}_{\mathrm{i}, \mathrm{a}} /\left(\mathbf{t}_{\mathbf{a}}{ }^{\prime} \cdot \mathbf{t}_{\mathbf{a}}\right)$

- Residuals e

How much structure/information remains after n-factors objects: $\Sigma \mathrm{e}^{2}$ over rows variables: $\Sigma \mathrm{e}^{2}$ over columns

Of the many diagnostics that become available from a PC analysis these five are the most used. Plotting them, keeping in mind the data its "secondary" structure mentioned before, forms the basis of exploratory data analysis.



In mathematical terminology the first factor is the least squares rank one approximation of $\mathbf{X}$. The second factor is the least squares approximation of all that remains (the residuals E).
Furthermore we want the scores and loadings (impose one and the other one comes for free) of successive factors to be orthogonal (inner product zero).

All that remains is to remove the so-called scaling ambiguity: if we divide the score-values by a constant c and multiply the loadings by this same $c$, we get the same solution. This holds for any $c$, thus we have infinitely many solutions! To solve this inconvenience we specify the length of loading vector $\mathbf{p}$ to one, thereby fixing c (but any other choice would have done).

Some terminology:
Bilinear modeling: PCA determines two least squares models, one in variable space and one in object space (also see NIPALS algorithm below).
Soft modeling: well... you impose the least squares criterion, but that is all. Not physical, chemical, psychological, whatever-al functions are predefined to get a particular solution.


The NIPALS algorithm (Non-linear Iterative PArtial Least Squares) is given to illustrate how to determine PCA factors. It consists of two least squares steps on the data table $\mathbf{X}$ : one to compute the loadings, one to compute the scores. These steps are repeated in sequence until no more changes in $\mathbf{t}$ and/or $\mathbf{p}$ are observed. In the middle we have to norm the loading $\mathbf{p}$ to get ride of the scaling ambiguity.



As we saw before: the first factor finds all the variation that the objects have in common, and goes through the origin. If there is an offset ("baseline") - like in the picture - this will often be the main variance. To get a more parsimonious model the data is almost always mean centered before analysis by subtracting the mean of a column from the entries in that column.

Other scaling-operations are sometimes used, where auto-scaling is the most famous one. In auto-scaling each variable get the same importance by setting the mean to zero (subtract the column average) and variance to one (divide by the column standard deviation).

Alternative transformations can lead to completely other model outcomes. The correct choice for transformations has to come from the problem domain (e.g. knowledge on the types of measurements preformed).


PCA can handle a modest amount of missing data. The model parameters ( $\mathbf{t}$ and $\mathbf{p}$ ) are estimated from a least squares fit by solving many simple equations. If there is a missing value in the table we just estimate the parameters from the equation that are complete. From the estimates we can than reconstruct the missing entry.

Whether this trick works depends on the amount of missing data and the structure in the matrix.

## Principal Component Analysis

how many factors: mathematical rank $v$. "true" rank


The important question in PCA-modeling is how many factors are relevant. The maximum mathematical rank of a matrix is the minimum of the number of rows and columns. But often the true rank of a system is much lower because the variables form (near) linear combinations (co-linearity). An important indicator is the percentage variance explained. This will always grow when increasing the model complexity. But if the increase from one factor to the next is small it is most likely noise you are fitting.

Another way to estimate the model complexity is cross validation. In PCA this is achieved by setting (non-overlapping) diagonals to missing. Next, the data table is reconstructed for an increasing number of factors, and the reconstructed diagonals are compared to the original entries. This procedure is repeated for all the diagonals. When the reconstruction error drops from one factor to the next, the extra component models real structure. When the reconstruction error flattens out, the factor describes primarily noise.


Ones you have a PCA model (= the correct number of loadings p) you can project new objects (= a series of new measurements $\mathbf{x}$ ) and get a new set of "principal" variables (= the score values $\mathbf{t}$ ).
In short: PCA defines coordinates/loadings $\mathbf{p}$ to filter the noise from the interesting directions and expresses the objects/samples in a new set of variables/scores $\mathbf{t}$.



Another interesting algorithm closely related to PCA is the Singular Value Decomposition (SVD). This method (frequently used inside PCA-software) decomposes the matrix $\mathbf{X}$ in three parts: the left singular vectors $\mathbf{U}$ (the object space), the right singular vectors $\mathbf{V}$ (the variable space) and the diagonal matrix $\mathbf{D}$ that has the singular values on it diagonal, sorted in decreasing order.

The diagonal matrix immediately shows the structure of the matrix: a few large singular values in the upper part represent the main phenomena, while the (near) zero diagonal elements below are the noise dimensions.

The left and right singular value matrices are both column-wise orthonormal, which has a nice property for matrix inversion we will use later on ( $\mathbf{U}^{-1}=\mathbf{U}^{\prime}$ ).

The SVD-decomposition can easily be turned into the PCA-solution. For two-dimensional data (so-called first order data) the two representations are equivalent. From SVD we also learn that PCA is an eigenvalue decomposition of the association matrix $\mathbf{X}^{\prime} . \mathbf{X}$ (or $\mathbf{X} . \mathbf{X}^{\prime}$ ).


So why not go between NIPALS (one factor at the time) and SVD (all factors at ones)? That is the Alternating Least Squares (ALS) algorithm, computing the bilinear models for the number of factors N you think are interesting. You start by picking a random $\mathbf{A}$ (the scores matrix) or $\mathbf{B}$ (the loading matrix), and iterate until you find the solution. The solution will span the same $N$-dimensional space as the first N-PCA factors (the least squares solution), but ...

... in this ALS we did not impose the additional criteria from PCA (orthogonal score/loading vectors, correct for scaling ambiguity), so our solution is in the same space but rotated.

This is however no problem. You are free to select a new base within the new factor space and any pair of axis will span the twodimensional plane (as long as they are not the same!).


The simple ALS algorithm opens the possibility to introduce (additional) constraints in a factor model. In this example the objectscores A are initialized as PCA scores $\mathbf{T}$. Then an iterative ALS algorithm is used to find a constraint factor model. The two constraints are: all entries in the variable-loadings larger than zero (nonnegativity), and fix the length of the loading vectors to one (eliminate scaling ambiguity). Constrains can assist in the interpretation of assumed physical phenomena in a data analysis (e.g. spectra in curve resolution problem).

By imposing the appropriate constrains (orthogonal solution vectors, sorted by size), the relation between ALS, NIPALS and SVD becomes obvious.

Notice that the two subspaces in this example (PCA versus ALS + constrains) are not the same. We chose to alter the model (and sacrifices fit) in favor of a non-negative loading vectors $\mathbf{B}$.

## Partial Least Squares regression

The next part of this document contains the transparencies used in the lectures on PLS, accompanied by some "telegram-style" notes on the main issues.

## Regression methods <br> Partial Least Squares (PLS)



- Find the relationship between predictor variables $(X)$ and response variables ( Y ) (independent/dependent)
- Aim is to predict unknown Y's for future X's (most often)
- Or, maybe to study $X$ under the effect of $Y$ (sometimes)


Partial Least Squares (PLS) can be used to build regression models between a predictor block $\mathbf{X}$ and a response vector $\mathbf{y}$. PLS has the possibility to form a regression model with a table $\mathbf{Y}$ (so-called PLS2), by determining a bilinear model for the response block and computing with the pseudo variables/scores $\mathbf{u}$. It is however not always favorable to do so, and since PLS is rather unique in this capability we will focus on the one response variable situation.


In its most simple form we are solving for $\mathbf{b}$, a vector of regression coefficients, and one coefficient for each variable in $\mathbf{X}$. If $\mathbf{y}$ is not mean centered an offset $b_{0}$ can included.


In a regression problem we are looking for a linear combination of the variables in $\mathbf{x}$ that can predict the response $\mathbf{y}$. This linear combination is defined by the regression vector $\mathbf{b}$, found from a least squares solution. This Multiple Linear Regression (MLR) solution can have some serious problems: is the predictor matrix $\mathbf{X}$ is (near) rank deficient the inverse is very hard to compute (unstable solution), and if the system is undetermined (more variables than objects) the solution is even undefined.


If the matrix $\mathbf{X}$ ("the system") is underdetermined (more variables/measurements than objects/samples) the inverse does not exist. If two or more variables are collinear the inverse of a overdetermined (more or equal number of observations than variables) might also not exist (singular or rank deficient system).

In practice many systems are "near rank deficient" - due to noise, rounding errors, etc. - but computing inverse will be very unreliable.


To overcome the problems in MLR we need to regularize the regression solution. One obvious candidate is to substitute the original variables by a smaller set of more relevant variables: the scores from a PCA analysis (the Principal Component Regression model). The score vectors are by definition orthogonal, leading to a well-conditioned inverse, and PCA can never extract more factors then the maximum rank of $\mathbf{X}$, so the inverse is never underdetermined.


The concept of a regularized inverse is best illustrated using the SVD to compute a PCR solution. We end up with computing the inverse of the diagonal matrix $\mathbf{D}$, which is easy because the inverse of a diagonal matrix is just the inverse of the elements on the diagonal.
Remember that $\mathbf{D}$ has the singular values (squared eigenvalues) in decreasing order as its entries. For the interesting part (upper corner) we have large singular values, and the inverse can be determined relatively easy. For the small singular values (lower corner; the noise) the inverse can "blow up": you end up with enormous numbers that completely dominate the inverse and thus the regression solution. If the last singular values are zero (something that seldom happens for real data) the inverse is undefined altogether.

| Regression methods Principal Component Regression - PCR |  |
| :---: | :---: |
| $\begin{aligned} & X=T \cdot P^{\prime}+E \quad \text { (pca) } \\ & y=T \cdot q+f \\ & T^{\prime} . y=T^{\prime} \cdot T . q \end{aligned}$ | $\mathrm{T}^{\prime} \cdot \mathrm{T}=\left(\begin{array}{lllll}10 & & & \\ & 1 & & \\ & 1 & & \\ & 0.01 & \\ & & & \sim 0\end{array}\right)$ |
| $\begin{aligned} & \left(T^{\prime} \cdot T\right)^{-1} \cdot T^{\prime} \cdot y=\left(T^{\prime} \cdot T\right)^{-1} \cdot\left(T^{\prime} \cdot T\right) \cdot q \\ & \left(T^{\prime} \cdot T\right)^{-1} \cdot T^{\prime} \cdot y=q \end{aligned}$ <br> (PCR solution, $y=X . b=X . P . q$ ) |  |



The observation on the other page is the heart of factor-models: you want to include the factors/principal components/latent variables that are significant but in doing so you disregard smaller dimensions, introducing a bias in the parameters/solution. The gain is more stable parameter estimation, because the difficulty in estimating factors in the smaller/unstable dimensions is avoided. Somewhere there is an optimal model complexity in number of factors: the bias-variance trade-off.

- Principal Component Analysis determines factors from a data table by making new 'pseudo'-variables (so-called principal components) from linear combinations of the original variables


## X

- The factors (principal components) are selected to explain as much information (variance) in the table as possible

By drawing the new axis/coordinates along the main (principal) direction in the data cloud


- The new variables eliminate redundant information (and filter noise) from the original data table


So far the new factor-coordinate system was selected to explain as much of the $\mathbf{X}$-variance as possible. There is however no guarantee that this set of loadings (or rather the scores formed by projection) is suitable for predicting $\mathbf{y}$.

In PLS an alternative set of coordinates in the form of loading-weights $\mathbf{W}$ is defined. The criteria of selection are to both explain variance in $\mathbf{X}$ (to get a stable inverse) and correlated the direction with $\mathbf{y}$ (or $\mathbf{u}$ in the case of multiple y's; to get good predictions): maximize the covariance. This more-or-less heuristic decision to make the two criteria equally important works well in practice and often leads to simpler models than e.g. PCR.


A (futile) attempt to visualize PLS: next to the three variable space of our X-block we have a one variable space of the $\mathbf{y}$-block (non- overlapping spaces!). The original factor 1 is tilted (out of the plane!) in the direction of the $y$-space. Hence, PLS is a skewed projection. It is easier to imagine this in the object space (three $\mathbf{x}$ - and one $\mathbf{y}$-variable both in eight dimensional spaces for this example) but that is impossible to draw.


The actual regression part inside PLS is a simple, univariate, linear regression between the predictor/X pseudo variables $\mathbf{t}$, and the responses/y or the pseudo variables $\mathbf{u}$ in case of multiple $y$ 's.

A point on regression factor-models: non-linearity in the parameters. A good way to detect non-linearity is to plot the two participants in the regression equation: $\mathbf{t}$ versus $\mathbf{y}$.

PLS is able handle modest non-linearity by simply including more factors. Many authors have attempted to tackle linearity issue by defining more complex regression equation inside the PLS algorithm (the so-called inner relation between $\mathbf{t}$ and $\mathbf{u} / \mathbf{y}$ ). Unfortunately the relations are often too complex (e.g. detector saturation) to find an easy/general solution.

A good tip when you expect non-linearity is to augment the data-block $\mathbf{X}$ with its squared form $\mathbf{X}^{2}$.


We use again the same small example as a reminder: socioeconomical status of 49 countries world-wide determined by 7 indicator variables/demographic.

In the regression problem we try to predict "percentage literate above age 14 " for the last 9 countries using the first 40 as training set.
R.F. Gunst and R.L. Mason "Regression Analysis and Its Application: A Data-Oriented Approach" Marcel Dekker, New York, p. 358(1980)


A three factor model gives slightly better $\mathrm{RMSP}_{c v}$ than a one factor model, but maybe only one factor is sufficient.

Regression coefficients show that "percentage literate above age 14" negatively related to physicians and infant deaths. This same response variable is positively related to predictor variable GNP. This is in good agreement with the findings in the Principal Component Analysis.


Overall the predictions of percentage literate from the other five socio-economical descriptors are good. Exceptional countries found in the PCA are also predicted relatively bad (Taiwan as a typical island state and USA which has (or had in 1980) a different university educational system).


The exceptional position of the USA also becomes clear by looking at the bi-plot for the first three PLS factors.


In factor regression problems we have to determine the optimal model complexity. This can be done by leave-one-object-out (or leave-a-particular-set-out) cross validation. We are looking for the minimum of the bias-variance RMSP-plot (in practice a no-longer-a-significantdrop, to get parsimonious models).

## Partial Least Squares Diagnostics

- Predictive performance

Cross validation/test-set

- Regression vector $b_{a}$
(Factor model, see PCA)
- Score values $t_{a}$ and $u_{a}$

Position of objects in the new factor-space

- Loading values $p_{\mathrm{a}}$ and $\mathrm{q}_{\mathrm{a}}$

Role of original variables in determining the new factor-space

- Percentage explained variance Closeness between factor space and original data spaces ( X or Y )
- Leverage $h_{a}$

How important is an object (or variable) compared to the rest of the data set

- Residuals $e_{a}$ and $f_{a}$

How much structurefinformation remains after n-factors

In PLS regression modeling we have the same + some more diagnostic tools as in PCA.

## Linear Algebra / PCA

In this part of the text we will use linear algebra methods and SVD/PCA to get some more experience in using the Matlab program. Let's start by generating some data. First we define two "true" loadings (Gaussian-curves of length 100) and two times ten "true" scores

```
>> p1 = gauss(100,65,8)';
>> p2 = gauss(100,35,6)';
>> t1 = [1:10]';
>> t2 = [3:-.1:2.1]' + randn(10,1)*0.5;
```

From this we generate our data matrix $X$, adding some noise, and plot the results

```
>> Xclean = t1*p1' + t2* p2';
>> X = Xclean + randn(size(t2* p2'))*0.003;
>> plot(1:100,Xclean,'b',1:100,X,'g');
>> grid;
```

We can define labels for the ten objects/samples in our set
>> ObjLab = ['\#0';'\#1';'\#2';'\#3';'\#4';'\#5'; '\#6';'\#7';'\#8'; '\#9']
To compute an SVD on this data table enter (function "diag" selects diagonal entries)
$\gg[U, D, V]=\operatorname{svd}(X)$;
$\gg d=\operatorname{diag}(D)$
$\gg \operatorname{plot}(1: 10, \operatorname{diag}(\mathrm{D}))$;
Check to see if the criteria defined in the figure on SVD are fulfilled: $U U^{\prime}=U^{\prime} U={ }_{10} I^{10}, V^{\prime}=V^{\prime} V={ }_{100} I^{100}, D=$ diagonal and sorted according to size, etc.

How much of the variation in $X$, in percentage, is explained by the first factor? We can compute this directly from the singular values in "D" $\gg d(1)^{\wedge} 2 / s u m(d . \wedge 2) * 100$

And how about the cumulative percentage variation explained for the first two factors? Try to dissect this Matlab command line $\gg \operatorname{sum}\left(d(1: 2) .^{\wedge} 2\right) / \operatorname{sum}(d . \wedge 2) * 100$

To get the first three PCA scores and loadings from the SVD decomposition simply compute $\gg \mathrm{T}=\mathrm{U}(:, 1: 3) * \mathrm{D}(1: 3,1: 3)$;
$\gg P=V(:, 1: 3)$;
$\gg \operatorname{plot}(1: 100, \mathrm{P}(:, 1: 3))$;
>> legend('P1','P2','P3')
Now try the same SVD decomposition for the clean data set (without noise, "Xclean" from above). What is the effect of adding noise? Try the Matlab commands "rank" and "cond" (= condition number).

To see e.g. the relation between the first true scores "t1" and the scores on the first PCA factor "T" ("text" puts the strings in "ObjLab" on position $x, y$ )
>> plot(t1,T(: , 1),'.');
>> xlabel('t1');
>> ylabel('T1')
>> text(t1,T(:,1),ObjLab)
The correlation coefficient between "t1" and the estimate "T(:,1)" can be computed as follows (on the diagonal is the correlation coefficient of a vectors with itself, which is of coarse always 1)
$\gg \mathrm{R}=\operatorname{corrcoef}(\mathrm{t} 1, \mathrm{~T}(:, 1))$
The squared correlation coefficient (= explained variance) is $\gg R(1,2) . \wedge 2$

How about the second set of true scores "t2"?

To study the alternating least squares (ALS) algorithm we need some more code (assign a random starting point in "A", "inv" computes the inverse of the matrix) $\gg \mathrm{A}=$ randn $(10,2)$;
$\gg$ for $a=1: 100, B=\left(\operatorname{inv}\left(A^{\prime} * A\right)^{*} A^{\prime *} X\right) ' ; A=X^{*} B^{*} \operatorname{inv}\left(B^{\prime *} B\right) ;$ end
This script does not specify a real convergence criterion. It just iterates/loops true the two steps one hundred times (the "for ... end" Matlab command), assuming the minimum in the least squares solution is found by then.

Lets plot the results
>> figure
$\gg \operatorname{subplot}(2,1,1) ; \operatorname{plot}(1: 100, \mathrm{~B}) ;$ title('loadings')
$\gg$ subplot(2,1,2); plot(t1,A(:,1),'.-',t2,A(:,2),'.-'); title('scores')

Try this ALS+plot procedure a couple of times (open a new figure window each time) with new random starting values in "A". What are the differences and similarities between solutions? How are these solutions different from the SVD/PCA results found earlier?

## Regression / PCR

In this part we will look at some regression equations in Matlab. Let's start by generating some data (same as in "PCA" part).
>> p1 = gauss(100,65,8)';
$\gg$ p2 = gauss(100,35,6)';
$\gg$ t1 $=[1: 10]$ ';
$\gg \mathrm{t} 2=[3:-.1: 2.1]^{\prime}+\operatorname{randn}(10,1) * 0.5 ;$
Then we generate our data matrices plus some noise
>> Xclean = t1*p1' + t2* p2';
$\gg X=$ Xclean + randn(size(t2*p2'))*0.003;
$\gg y=t 1+t 2+\operatorname{randn}(\operatorname{size}(t 2))^{*} 0.1 ;$
>> ObjLab = ['\#0';'\#1';'\#2';'\#3';'\#4';'\#5'; '\#6'; '\#7'; '\#8'; '\#9']
To compute an SVD on this data enter
$\gg[U, D, V]=\operatorname{svd}(X)$;
From this decomposition we compute the Multiple Linear Regression (MLR) solution ( 10 object by 100 variables, so the maximum number of factors is 10) and plot it
$\gg$ b10 = V(:, 1:10)*inv(D(1:10,1:10))*U(:, 1:10)'*y
>> plot(1:100,b10);
This will look like a lot of noise! Why is that? We included too many factors. Let's look at the 1, 2 and 3-factor Principal Component Regression solutions instead

```
>> b1 = V(:,1:1)*inv(D(1:1,1:1))*U(:,1:1)'*y
>> b2 = V(:,1:2)*inv(D(1:2,1:2))*U(:,1:2)'*y
>> b3 = V(:,1:3)*inv(D(1:3,1:3))*U(:,1:3)'*y
>> plot(1:100,[b1 b2 b3]);
>> legend('b1','b2','b3')
```

We can also predict (or rather "fit" in this case) the $y$-values
>> ypl = $X^{*} b 1$;
$\gg y p 2=X^{*} b 2 ;$
$\gg y p 3=X * b 3 ;$
>> plot(y,yp1,'.-', y,yp2,'.-', y,yp3,'.-'); grid
>> legend('1 factor','2 factors','3 factors');
>> line([2 14],[2 14])
The "line" command plots a line: offset 0 , slope $45^{\circ}$.

## NIR with temperature effects

A set of Near InfraRed (NIR) ${ }^{*}$ spectra is available in the Matlab data-file "NIRdata.mat". This set (three X-blocks "specXX": 19 x 512; spectra recorded at $30^{\circ} \mathrm{C}, 50^{\circ} \mathrm{C}$ and $70^{\circ} \mathrm{C}$ ) is part of a larger triangular design formed by mixtures of ethanol (first column in "conc": the $Y$ block with mass fractions of the constituents), water (second column) and iso-propanol (third column). The corner points of the design (the pure components/spectra) are removed from the main set, but are available in the data-file. The Matlab variable "ObjLab" contains appropriate name-labels for the objects (see first figure below for the design).

[^0]To plot the mixture design points:
>> clear all
$\gg$ close all
$\gg$ load NIRdata
>> plotNIRDesign(conc)
$\gg$ text(conc(:, 1$),$ conc(:, 2$),$ conc(: , 3), ObjLab)
>> xlabel('ethanol')
>> ylabel('water')
>> zlabel('iso-propanol')


Use the rotate button to get a better view of the design!
Next, have a look at the temperature influence by plotting the pure components of the mixtures:

```
>> figure
>> subplot(3,1,1)
>> plot(wavelen,[ethanol30; ethanol50; ethanol70])
>> legend('30C','50C','70C',2)
>> title('ethanol')
>> ylabel('A.U.')
>> grid
>> subplot(3,1,2)
>> plot(wavelen,[water30; water50; water70])
>> title('water')
>> ylabel('A.U.')
>> grid
>> subplot(3,1,3)
```

>> plot(wavelen,[isopropanol30; isopropanol50; isopropanol70])
>> xlabel('wavelength (nm)')
$\gg$ ylabel('A.U.')
$\gg$ grid
>> title('iso-propanol')


Use the zoom button for a detailed view. What is the chemical nature of the changes?

The different mixture-spectra look as follows:
>> figure
$\gg$ subplot $(3,1,1)$
>> plot(wavelen,spec30,'b')
>> ylabel('A.U.')
$\gg$ grid
$\gg$ subplot $(3,1,2)$
>> plot(wavelen,spec50, 'g')
>> ylabel('A.U.')
$\gg$ grid
>> subplot $(3,1,3)$
>> plot(wavelen,spec70,'r')
>> xlabel('wavelength (nm)')
>> ylabel('A.U.')
$\gg$ grid


A SVD/PCA analysis of the $30^{\circ} \mathrm{C}$ data-set could go as follows:
>> [U30,D30,P30] = svd(spec30);
>> T30 = U30*D30;
Have a look at the loading-vectors (pseudo spectra); how do they compare to the raw spectra? How much variance is explained by the first four factors (use D30)? What is the rank of this system? What
should the rank from a chemical point of view? And what about closure?

Can you still find the design (use rotate) plotting the scores of the first three factors?
>> figure
>> plotNIRDesign(T30(:,1:3))
>> xlabel('PC1'); ylabel('PC2'); zlabel('PC3');
If you rotate this plot you will see it looks like a plane (representing the mixture design) in a 3D room. What we did not do in our SVD/PCA analysis above is mean centering the data before decomposition. But this is by now "easy" in Matlab (hints: first we determine the size of the matrix, then we compute the mean spectrum - thus over the columns in the row direction; finally a trick: we subtract a matrix from the data build from the outer product of a vector of ones and the mean spectrum. What happens here? Check the outer product.)

```
>> [N,M] = size(spec30);
>> mean spec30 = mean(spec30,1);
>> spec30c = spec30 - ones(N,1)* mean_spec30;
```

Check the result of the mean centering pre-processing
>> figure
>> subplot( $2,1,1$ );
>> plot(wavelen, spec30, 'b',wavelen, mean_spec30, 'g'); grid;
>> subplot(2,1,2); plot(wavelen,spec30c, 'r'); grid;
And try SVD/PCA on this
>> [U30c,D30c,P30c] = svd(spec30c);
>> T30c = U30c*D30c;
>> figure
>> plotNIRDesign(T30c(:,1:3))
>> xlabel('PC1'); ylabel('PC2'); zlabel('PC3');
What happend to the rank and the singular values in "D30c"? What happened to the image of the design in the score plot?

Try the same for the mixture-sets recorded at $50^{\circ} \mathrm{C}$ and $70^{\circ} \mathrm{C}$. What effect does temperature have on the results?

## Basic Statistics

To study some basic we will use a small example of 10 pH measurement from one solution. The data is as follows
>> clear all
>> close all
$\gg \mathrm{pH} 1=[4.905 .065 .055 .175 .064 .945 .044 .905 .005 .00$ ]';
The first thing to do when collecting new data is to plot it!!!
In this case we have the 10 observations, a so-called box-andwhiskers plot, the histogram (with a normal or Gaussian distribution imposed) and the normal probability plot

```
>> figure
>>subplot(2,2,1);
>> plot(pH1,ones(1,length(pH1)), '.','MarkerSize',20);
>> grid; title('pH1');
>> subplot(2,2,2);
>> boxplot(pH1,'notch','on');
>> title('Box plot pH1');
>> subplot(2,2,3);
>> normplot(pH1);
>> title('Normal prob. plot pH1');
>> subplot(2,2,4);
>> histfit(pH1); grid;
>> title('Histogram pH1')
```



This data-set is a little small to study the normal probability plot and the histogram function. Try the following code which generates to vectors with 200 random numbers each
$\gg$ figure
$\gg x n=\operatorname{randn}(1,200) ; x=\operatorname{rand}(1,200)$;
$\gg$ subplot( $2,2,1$ ); histfit( $x n$ );
$\gg$ subplot $(2,2,2)$; histfit(x);
$\gg$ subplot $(2,2,3)$; normplot( $x n$ );
$\gg$ subplot( $2,2,4$ ); normplot(x);
What is the difference between random numbers generated by "rand" and "randn" (use the help function)?

Compute the important statistics for this pH data-set: mean standard deviation, relative standard deviation and standard error of the mean
(make sure you understand both the meaning of these numbers AND the Matlab commands to calculate them)

```
>> m_pH1 = mean(pH1)
>> s_pH1 = std(pH1)
>> RSD_pH1 = s_pH1/m_pH1*100
>> SE_m_pH1 = std(pH1)/sqrt(10)
```

The $95 \%$ confidence interval can be formulated as follows ("tinv" is a function equivalent to the student-t lookup table; note $5 \%$ two-sided it $2.5 \%$ or 0.25 for each side)

```
>> t95p = tinv(0.975,length(pH1)-1)
>> L = m_pH1 - t95p*SE_m_pH1;
>> U = m_pH1 + t95p*SE_m_pH1;
>> disp([num2str(L,3)' < mu pH1 < ' num2str(U,3)])
```

Next step is the comparison of two data-sets. We make a second vector with pH values and plot the result

```
>> pH2 = [5.10 5.07 5.21 4.91 5.14 5.19 5.17 5.16 5.10 5.17]';
>> figure
>> subplot(2,1,1);
>> plot(pH1,ones(1,length(pH1)),'.b', pH2,2*ones(1,length(pH2)),'.g','MarkerSize',20);
>> grid; title('pH1 and pH2');
>> subplot(2,1,2);
>> boxplot([pH1 pH2],'notch','on'); title('Box plot pH1 and pH2');
```

Are the two pH -series the same? To compare them we can use the socalled Fisher F-test, based on the ration between Treatments (= the two pH series) and the total error/uncertainy ("finv" is a function equivalent to the F lookup table; again, make sure you understand both the meaning of these numbers and the Matlab commands to calculate them)

```
>> MStreatment = sum(10*([mean(pH1);mean(pH2)]-mean([pH1; pH2])).^2)/(2-1)
>> MSerror = sum([(pH1-mean(pH1)).^2;(pH2-mean(pH2)).^2])/(20-2)
>> F = MStreatment/MSerror
>> c = [0.75 0.90 0.95 0.975 0.99];
>> [100-c* 100; finv(c,1,18)]
```

From the list we see that finding these $2 \times 10$ numbers, under the assumption that the $\mathrm{pH}^{\prime} \mathrm{s}$ of the two series where the same, is smaller then 1\%.

## Linear Regression

We will use the NIR spectra from above to implement univariate linear regression in Matlab. If we want to predict the water contents in our samples we could use e.g. the absorbance value at 950 nm (confirm this from the spectral plots above; closely study the way we find 950 nm in the variable "wavelen"; remember water is the second column in "conc")

```
>> clear all
>> close all
>> load NIRdata
>> y = conc(:,2);
>> [dummy,index950] = min(abs(wavelen-950));
>> wavelen(index950)
>> X = spec30(:,index950);
```



A plot of water fraction as a function of the absorbance at 950 nm looks as follows

```
>> figure
>> plot(X,y,'.')
>> grid; xlabel('Abs. 950nm'); ylabel('Water (fraction) reference');
```

Next we form our matrix of independent $\mathbf{X}$, where we add a column of ones for the offset (the $b_{0}$ term). From this matrix and the dependent $y$ we can estimate our unknowns in the regression equation $y_{\text {water }}=b_{0}$ $+X_{950 n m} \cdot b_{1}$;

```
>> X = [ones(size(X)) X];
```

$\gg b=\operatorname{inv}\left(X^{\prime} * X\right) * X^{\prime *} y$

To see how our model performs we estimate our water fractions and make the regular predicted-versus-reference plot (red line is the ideal line: offset zero, slope 1), the Root Mean Square error of Prediction for fit (Why is this fit, and nor real prediction? Check the formula for RMSP, working from the inside to the outside), and the correlation coefficient

```
\(\gg\) yhat \(=X^{*}\) b;
\(\gg\) figure
>> plot(y,yhat, '.',[[0 1],[0 1],'r')
>> grid; axis square; xlabel('Water (fraction) reference');
>> ylabel('Water (fraction) predicted/fitted');
\(\gg\) rmsp_fit \(=\) sqrt(mean((y-yhat).^2))
>> corrcoef(y,yhat)
```

Alternatively we can plot our estimates together with the original absorbances

```
>> figure
```

>> plot(X(:,2),y,'.',X(:,2),yhat,'+g',[0 max(X(:,2))],[[1 0]*b [1 max(X(:,2))]*b],'b')
>> grid; xlabel('Abs. 950nm'); ylabel('Water (fraction) reference');

Why is this regression model of NIR spectra to predict water not very good (hint: have another look at the spectra)? And how could we improve it (hint: what about switching to MLR or PLS including e.g. absorbance wavelength 908 nm - now we are back at the origin of Chemometrics again)?

Here is some "quick and dirty" steps to achieve MLR $\gg$ [dummy, index908] $=\min ($ abs(wavelen-908));
$\gg X=[X \operatorname{spec} 30(:$, index908)];
$\gg$ b_mlr $=\operatorname{inv}\left(X^{\prime *} X\right) * X^{\prime *} y$
$\gg$ yhat_mlr $=X^{*}$ b_mlr;

## Linear Algebra

## Linear Algebra

$\mathbf{C}$ is of size $n \times m=3 \times 2$ ( 3 rows, 2 columns), $\mathbf{a}$ and $\mathbf{b}$ are of size $3 \times 1$


Conventions:

- matrix = build from vectors
- indexing elements in matrices (= scalars): first row index, then column index







## Linear Algebra <br> Singular Value Decomposition (SVD)

We set out to solve for the unknowns in a linear system of equations
$\left.\left[\begin{array}{ccc}2 & 1 & 1 \\ 4 & 1 & 0 \\ -2 & 2 & 1\end{array}\right] \begin{array}{l}b_{1} \\ b_{2} \\ b_{3}\end{array}\right]=\left[\begin{array}{c}1 \\ -2 \\ 7\end{array}\right]$
$\mathbf{X b}=\mathbf{y} \quad \mathbf{X}^{-1} \mathbf{X b}=\mathbf{X}^{-1} \mathbf{y} \quad \mathbf{b}=\mathbf{X}^{-1} \mathbf{y}$
$\mathbf{X}^{-1} \mathbf{X}=\left[\begin{array}{ccc}1 / 8 & 1 / 8 & -1 / 8 \\ -1 / 2 & 1 / 2 & 1 / 2 \\ 5 / 4 & -3 / 4 & -1 / 4\end{array}\right]\left[\begin{array}{ccc}2 & 1 & 1 \\ 4 & 1 & 0 \\ -2 & 2 & 1\end{array}\right]=\left[\begin{array}{lll}1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1\end{array}\right]$
$\mathbf{b}=\mathbf{X y}=\left[\begin{array}{ccc}1 / 8 & 1 / 8 & -1 / 8 \\ -1 / 2 & 1 / 2 & 1 / 2 \\ 5 / 4 & -3 / 4 & -1 / 4\end{array}\right]\left[\begin{array}{c}1 \\ -2 \\ 7\end{array}\right]=\left[\begin{array}{c}-1 \\ 2 \\ 1\end{array}\right]=\mathbf{b}$

Basic Statistics












## Regression model <br> $\mathrm{y}=\mathrm{X} . \mathrm{b}$

- We will limit ourselves to a so-called linear, additive models of the form:

$$
y=b_{0}+b_{1} \cdot x_{A}+b_{2} \cdot x_{B}+b_{3} \cdot x_{A} \cdot x_{B}+e
$$

Explanation:

- $y$ is dependent, $x_{A}$ and $x_{B}$ are the independents
- $b_{0}$ is the offset or grant average
- $b_{1}$ and $b_{2}$ are the weight factors for the main effects $A$ and $B$, respectively
- $b_{3}$ shows the importance for the interaction effect
- $e$ is the error, everything not modeled by the function
- linear (better: linear in the parameters): b.x, b. $x^{2}$, b.x.y, etc.; but not $x^{b}$, $e^{b x}$
- additive: total effect $y$ is sum of separate contributions


## Regression model <br> $\mathrm{y}=\mathrm{X} . \mathrm{b}$

- To find the b's we have to solve the Multiple Linear Regression (MLR) model:
$y=X . b \quad X^{\prime} \cdot y=X^{\prime} . X \cdot b \quad \rightarrow \quad\left(X^{\prime} \cdot X\right)^{-1} \cdot X^{\prime} y=\left(X^{\prime} \cdot X\right)^{-1} \cdot\left(X^{\prime} X\right) \cdot b$

$\rightarrow \quad \mathbf{b}=\left(\mathbf{X}^{\prime} \cdot \mathbf{X}\right)^{-1} \cdot \mathbf{X}^{\prime} \cdot \mathbf{y}$



## Linear regression <br> Concept of least squares

Simple regression model $y(i)=b_{0}+b_{1} \cdot x(i)+e(i)$

Assume independent variables
$y(1), y(2) \ldots y(n)$ :
Mean $=b_{0}+b_{1} \cdot x(i)$
Variance $=\sigma_{y}^{2}$

Assume independent errors $e(1), e(2) \ldots e(n)$
Mean = 0
Variance $=\sigma^{2}{ }_{y}$


Linear Regression

Slope b1 = change in $y$ that accompanies a unit change in $x$

Offset $\mathrm{b}_{0}=$ zero-response (non-response, 'bias')
Least squares = minimize the squared sum of the residuals

## Linear regression <br> Concept of least squares

Simple regression model $y(i)=b_{0}+b_{1} \cdot x(i)+e(i)$
True (unknown) model $y(i)=\beta_{0}+\beta_{1} \cdot x(i)$
Statistical errors e(i) in the observation

Estimates $b_{1}$ 's, (app.) normally distributed:

$$
E\left(b_{1}\right)=\beta_{1} \text { and } S E\left(b_{1}\right) \approx \frac{1}{\sqrt{n}} x \frac{s_{y}}{s_{x}}
$$

Better estimate for $b_{1}$ :

- Increase sample size $n$
- Reduce s.d. y(i) $\rightarrow$ e(i)
- Increase s.d. x(i) (leverage)







[^0]:    Florian Wülfert, Wim. Th. Kok and Age K. Smilde "Influence of temperature on vibrational spectra and consequences for the predictive ability of multivariate models" Anal.Chem.(1998)1761-1767

