Statistical and numerical methods of data analysis II

V. Blobel

Universität Hamburg

Desy, **June** 21 (+14, 28), **2001**

- Parameter estimation and information
- Least squares methods

Parameter estimation and information

- Estimates
- Samples
- Information
- Method of moments
- Robust estimates

Estimates

Given a sample $x_1, \ldots, x_i, \ldots, x_n$ (or short $\{\boldsymbol{x}\}$) of measured values from a distribution $p(x; a)$.

Parameter estimate $\quad \widehat{a} = \widehat{a}(x_1, x_2 \ldots x_n)$

is a function of the random variables $x_1, \ldots, x_i, \ldots, x_n$ and is itself a random variable.

Example $\quad \widehat{a} = \dfrac{1}{n}\sum_{i=1}^{n} x_i \quad$ mean value of x from $p(x; a)$

Properties of estimate $\quad E[\widehat{a}] = E[\widehat{a}(x_1, x_2 \ldots x_n)]$

Variance $\quad V[\widehat{a}] = V[\widehat{a}(x_1, x_2 \ldots x_n)]$

Information

Given a sample $x_1, \ldots, x_i, \ldots, x_n$ (or short $\{\boldsymbol{x}\}$) of measured values from a distribution $p(x; a)$ (i.e. normalized density).

What is the information in the sample about the parameter(s) a?

What is the accuracy that can be achieved from the sample about the parameter(s) a?

The likelihood function as *joint density* of the observed values of the random variable x:

Likelihood function $\quad \mathcal{L} = \prod_{i=1}^{n} p(x_i; a)$

with normalization

$$\int_{\Omega} \mathcal{L}\, \mathrm{d}x_1 \mathrm{d}x_2 \ldots \mathrm{d}x_n = 1$$

Information inequality

$$\text{Information} \quad I(a) = E\left[\left(\frac{\partial \ln \mathcal{L}}{\partial a}\right)^2\right] = \int_\Omega \left(\frac{\partial \ln \mathcal{L}}{\partial a}\right)^2 \mathcal{L}\, \mathrm{d}x_1 \mathrm{d}x_2 \ldots \mathrm{d}x_n$$

This is the definition of *information*, where $\mathcal{L}$ is the joint density of the n observed values of the random variable x.

$$\text{Information inequality} \qquad V[\widehat{a}] \geq \frac{1}{I}$$

The *inverse* of the information $I_n(a)$, or short I, is the lower limit of the variance of the parameter estimate $\widehat{a}$.

The inequality is also called Rao-Cramér-Frechet inequality, and is valid in this form for any unbiased estimate $\widehat{a} = \widehat{a}(x)$.

Efficiency and bias

Definition of the *efficiency* of an estimator:

For an unbiased estimator $\hat{a}$ one can define the *efficiency* $\mathrm{eff}(\widehat{a})$ by the ratio of the mininal to the actual variance:

$$\text{Efficiency} \qquad \mathrm{eff}(\widehat{a}) = \frac{I^{-1}}{V[\widehat{a}]} \qquad 0 \leq \mathrm{eff}(\widehat{a}) \leq 1$$

The actual efficiency of an estimator depends on the specific problem and method.

Variance limit in case of a bias $B_n(\widehat{a}) = E[\widehat{a}] - a_{\text{true}} \neq 0$:

$$\text{Variance} \qquad V[\widehat{a}] \geq \frac{(1 + \partial B/\partial a)^2}{I}$$

Case of several variables

Case of m variables $a_1, \ldots, a_j, \ldots, a_m$: information I becomes a m-by-m symmetric matrix $\boldsymbol{I}$ with elements

$$I_{jk} = E\left[\frac{\partial \ln \mathcal{L}}{\partial a_j}\frac{\partial \ln \mathcal{L}}{\partial a_k}\right] = -E\left[\frac{\partial^2 \ln \mathcal{L}}{\partial a_j \partial a_k}\right]$$

The minimal variance $\boldsymbol{V}[\hat{\boldsymbol{a}}]$ of an estimate $\hat{\boldsymbol{a}}$ is given by the inverse of the information matrix $\boldsymbol{I}$:

minimal variance	$\boldsymbol{V}[\hat{\boldsymbol{a}}] = \boldsymbol{I}^{-1}$

Example: Weighing four object with a balance

Weighing four objects $a_j, j = 1, \ldots 4$

Direct approach: weigh each object separately in four independent measurements.

Alternatively: try to improve the precision of four measurements by weighing suitable combinations of the objects:

$$x_i = \sum_{j=1}^{4} A_{ij} a_j \qquad \text{with} \qquad A_{ij} = -1,\ 0 \quad \text{or} \quad +1$$

Assumptions:

- measurements follow a normal distribution law, and
- standard deviation σ in each measurement is always the same.

Example: Weighing four object with a balance (2)

Likelihood function

$$\mathcal{L}(\boldsymbol{a}) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^4 \exp\left[-\frac{1}{2\sigma^2}\sum_{i=1}^{4}\left(x_i - \sum_{j=1}^{4} A_{ij}a_j\right)^2\right]$$

in a combination of measured quantities $x_i, i = 1, \ldots 4$ with $\boldsymbol{x} = \boldsymbol{A}\boldsymbol{a}$.

The elements of the information matrix are given by

$$I_{jk} = E\left[-\frac{\partial^2 \ln \mathcal{L}}{\partial a_j \partial a_k}\right] = \frac{1}{\sigma^2}\sum_{i=1}^{4} A_{ij}A_{ik} \qquad \text{or short} \quad \boldsymbol{I} = \frac{1}{\sigma^2}\boldsymbol{A}^T\boldsymbol{A}$$

Example: Weighing four object with a balance (3)

Direct approach: matrix $\boldsymbol{A}$ is the unit matrix ($A_{ij} = \delta_{ij}$)

$$\boldsymbol{I} = 1/\sigma^2\,\mathbf{1} \qquad \boldsymbol{V}\,[\widehat{a}] = \sigma^2\,\mathbf{1}$$

i.e. the objects a_j are measured with standard deviation σ.

Alternatively: linearly independent combination given by the matrix

$$\boldsymbol{A} = \begin{pmatrix} 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad \boldsymbol{I} = \frac{4}{\sigma^2}\mathbf{1} \qquad \boldsymbol{V}\,[\widehat{a}] = \frac{\sigma^2}{4}\mathbf{1}$$

Alternative method is preferable: objects a_j are now measured with standard deviation $\sigma/2$.

In both cases the measured values are uncorrelated.

Experimental design – Example

Decay angular distribution of an elementary particle

The decay angular distribution is given by

$$p(x|a) = \frac{1}{2}(1 + ax) \qquad \text{with} \quad x \equiv \cos\vartheta$$

The value of the parameter a is expected to be around $a = -1/3$. The question is, how many decays have to be observed to determine the parameter a with a relative precision of 1 %.

The information is calculated for $n = 1$:

$$I_1 = \int_{-1}^{+1} \frac{1}{p(x,a)} \left(\frac{\partial p}{\partial a}\right)^2 \mathrm{d}x = \begin{cases} \frac{1}{2a^3}\left[\ln\frac{1+a}{1-a} - 2a\right] & \text{if} \quad a \neq 0 \\ \frac{1}{3} & \text{if} \quad a = 0 \end{cases}$$

The standard deviation depends on the value of the parameter. For $a = -1/3$ the value is $\sigma = \sqrt{2.797/n}$ and it follows that more than 2.52×10^5 decays have to be measured in order to obtain a standard deviation $\sigma < 0.01\ a$.

Method of moments

The expectation of a function $h(x)$ of the random variable, given the density $f(x;a)$, is itself a function $g(a)$ of the parameter a:

$$E\left[h(x)\right] = \int_{-\infty}^{+\infty} h(x)\, f(x,a)\, \mathrm{d}x = g(a)$$

Compared to the average value of the function $h(x)$ for the data sample:

$$\overline{h(x)} = \frac{1}{n}\sum_{i=1}^{n} h(x_i)\ .$$

Principle of the method of moments: equate the theoretical expectation value and the empirical average of the function $h(x)$, to obtain the parameter estimate $\widehat{a}$.

Method of moments contnd.

$$g(\widehat{a}) = \overline{h(x)} \qquad \widehat{a} = g^{-1}\left(\overline{h(x)}\right)$$

which requires the function g^{-1} inverse to $g(a)$. According to the central limit theorem the mean value $\overline{h(x)}$ is asymptotically normal distributed with variance

$$V\left[\overline{h(x)}\right] = \frac{1}{n} V\left[h(x)\right]$$

The variance of the estimate is calculated using standard error propagation.

The approximate normality of the estimate is often destroyed by the inverse function $g^{-1}\left(\overline{h(x)}\right)$. The estimate has a potential bias $B_n(\widehat{a}) \propto 1/n$.

Robust estimates

Many estimates are based on the normal distribution. The normal distribution is the most important distribution in all statistics.

> Under random sampling, as the sample size gets large, the limiting distribution of the sample mean is normal." (Central limit theorem)

> "Everyone believes in the normal law of errors, the experimenters because they think it is a mathematical theorem, the mathematicians because they think it is an experimental fact" (Poincaré)

In the normal distribution only two numbers are required to tell us everything about the probabilities associated with a random variable:

the mean μ and the variance σ^2

Departures from normality

But: **Distributions are never normal!**

Departures from normality or from any other assumed distrinution are rather common in practice. Random samples ...

- often have outliers: unusual large or small values among a sample,
- can have heavy tails (thicker than a normal distribution),
- can be highly skewed.

Estimates of statistical parameters can be sensitively influenced by such deviations.

> Robustness signifies insensivity to small deviations from the assumptions.

Example: contaminated normal distribution

A large, randomly mixed batch of n "good" and "bad" observations of the same quantity μ is given.

- A "good" observation x_i from $N(\mu, \sigma^2)$ occurs with probability $(1-\varepsilon)$ (for small ε);
- a "bad" observation x_i from $N(\mu, K^2\sigma^2)$ occurs with probability ε with e.g. $K = 3$.

i.e. the x_i are independent, identically distributed with the underlying distribution

$$f(x) = (1-\varepsilon)\cdot n\left(\frac{x-\mu}{\sigma}\right) + \varepsilon \cdot n\left(\frac{x-\mu}{K\sigma}\right) \qquad \text{(contaminated normal)}$$

($n(x)$ = standard normal density).

How should the "scatter" (standard deviation) around the mean be measured?

Estimates for the scatter around the mean

Possible formulas for scatter:

$$s_n = \left[\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^2\right]^{1/2} \qquad \text{mean square deviation}$$

$$d_n = \frac{1}{n}\sum_{i=1}^{n}|x_i-\bar{x}| \qquad \text{mean absolute deviation}$$

Conventional answer: for normal observations s_n is about 12 % more efficient than d_n.

$$\lim_{n\to\infty} s_n = \sigma \qquad \lim_{n\to\infty} d_n = \sqrt{2/\pi}\,\sigma \approx 0.80\,\sigma$$

Efficiency for contaminated normal

But the distribution is not normal, it is the contaminated normal distribution!

Asymptotic relative efficiency of d_n relative to s_n (for $K = 3$):

$$A(\varepsilon) = \lim_{n\to\infty}\frac{V[s_n]/\left(E[s_n]\right)^2}{V[d_n]/\left(E[d_n]\right)^2} = \frac{1}{4}\frac{\dfrac{3(1+80\varepsilon)}{(1+8\varepsilon)^2}-1}{\dfrac{\pi(1+8\varepsilon)}{2(1+2\varepsilon}-1}$$

ε	0	0.001	0.002	0.005	0.01	0.02	0.05	0.10	0.25
$A(\varepsilon)$	0.876	0.948	1.016	1.198	1.439	1.752	2.035	1.903	1.371

Just 2 bad observations in 1000 are sufficient to get a better efficiency for the mean absolute deviation, and for $\varepsilon = 5\%$ the efficiency is better by a factor of > 2.

Method for mean value I: Trimmed mean

The random sample is x_1, $x_2, \ldots x_i \ldots x_n$. The observations are written in ascending order: $x_{(1)}$, $x_{(2)}, \ldots x_{(i)} \ldots x_{(n)}$.

A value of γ must be chosen, for example $\gamma = 0.2$

g = value of γn rounded down to the nearest integer.

The sample trimmed mean is computed by removing the g smallest and g largest observations:

$$\hat{x}_t = \frac{x_{(g+1)} + \cdots + x_{(n-g)}}{n - 2g}$$

Method for mean value II: Winsorized mean

Winzorization of a random sample x_i: the g smallest values are pulled in and set equal to x_{g+1} and the g largest values are pulled in and set equal to x_{n-g}

$$z_i = \begin{cases} x_{g+1} & \text{if} \quad x_i \leq x_{g+1} \\ x_i & \text{if} \quad x_{g+1} < x_i < x_{n-g} \\ x_{n-g} & \text{if} \quad x_i \geq < x_{n-g} \end{cases} \qquad \bar{x}_w = \frac{1}{n} \sum_{i=1}^{n} z_i$$

Another method is discussed in the Least squares chapter.

Least squares methods

- Systems of linear equations – singular value decomposition
- Linear least squares – properties of least squares estimates
- Normal equations for uncorrelated data
- Inversion by partitioning
- Orthogonal polynomials
- Least squares with constraints
- Alignment problems

System of linear equations

Solving the matrix equation

$$\boldsymbol{A}\boldsymbol{a} = \boldsymbol{y} \qquad \boldsymbol{A}\boldsymbol{a} \cong \boldsymbol{y} \qquad m \times n \text{ matrix } \boldsymbol{A}$$

for an n vector $\boldsymbol{a}$ of unknowns (parameters).

Distinguish cases:

- $n = m$ exact solution usually exists
- $m > n$ overdetermined system – obtain best "approximate" solution
- $m < n$ underdetermined system – insufficient data

$$\boldsymbol{A}\cdot\boldsymbol{a}=\boldsymbol{y} \qquad\qquad \boldsymbol{A}\cdot\boldsymbol{a}\cong\boldsymbol{y}$$

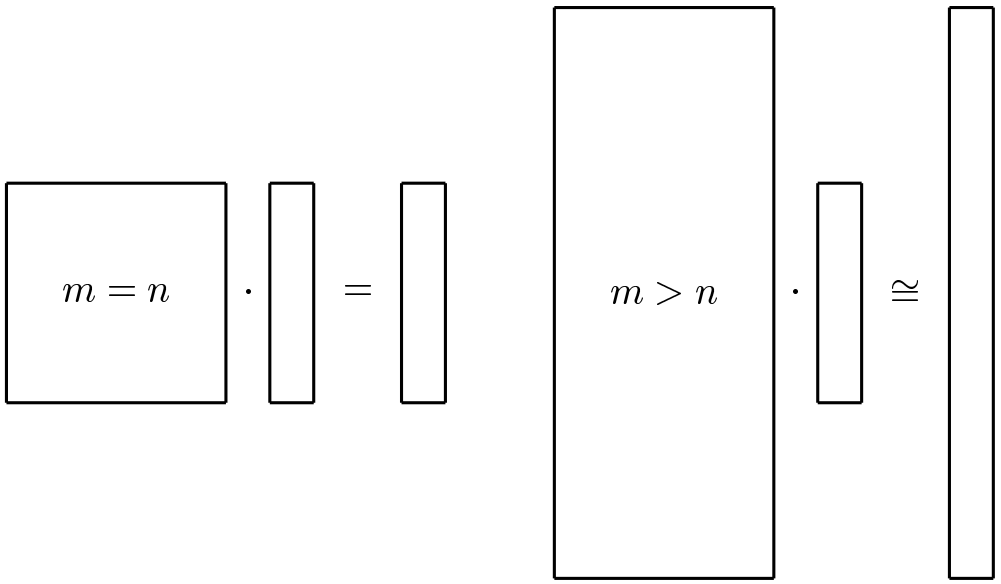

Overdetermined system ($m > n$)

In general no solution vector $\boldsymbol{a}$ exist, i.e. vector of residua $\boldsymbol{r}$, defined by

$$\boldsymbol{r}=\boldsymbol{A}\boldsymbol{a}-\boldsymbol{y}\,,$$

never zero.

What is the "best" solution?

Minimum value of the square of the vector of residua:

$$\begin{aligned}\boldsymbol{r}^T\boldsymbol{r} &= (\boldsymbol{A}\boldsymbol{a}-\boldsymbol{y})^T(\boldsymbol{A}\boldsymbol{a}-\boldsymbol{y})\\ &= \boldsymbol{a}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{a}-2\boldsymbol{a}^T\boldsymbol{A}^T\boldsymbol{y}+\boldsymbol{y}^T\boldsymbol{y} \;=\; \text{Minimum}\,.\end{aligned}$$

The minimum is determined, by setting the derivative of the above expression to zero:

$$\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{a}-\boldsymbol{A}^T\boldsymbol{y}=0 \quad \text{oder} \quad \boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{a}-\boldsymbol{y})=\boldsymbol{A}^T\boldsymbol{r}=0\,.$$

These equation are called *normal equations*; the vector $\boldsymbol{r}$ of residua is *normal* with respect to the column vectors of matrix $\boldsymbol{A}$.

Solution

Often the system has a unique solution

$$\boldsymbol{a} = \left(\boldsymbol{A}^T\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{y}$$

The symmetric matrix $\boldsymbol{C} = \boldsymbol{A}^T\boldsymbol{A}$ and the vector $\boldsymbol{c} = \boldsymbol{A}^T\boldsymbol{y}$ are formed:

$$\boxed{\boldsymbol{C}} = \boxed{\boldsymbol{A}^T} \times \boxed{\boldsymbol{A}} \qquad \boxed{\boldsymbol{c}} = \boxed{\boldsymbol{A}^T} \times \boxed{\boldsymbol{y}}$$

The solution vector $\boldsymbol{a}$ is determined by the solution of the matrix equation $\boldsymbol{C}\boldsymbol{a} = \boldsymbol{c}$:

$$\boxed{\boldsymbol{a}} = \boxed{\boldsymbol{C}^{-1}} \times \boxed{\boldsymbol{c}}$$

A solution in different form ...

The symmetric matrix $\boldsymbol{C}$ can be diagonalized:

$$\boxed{\boldsymbol{C}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{D}} \times \boxed{\boldsymbol{U}^T}$$

Matrix $\boldsymbol{U}$ is orthonormal and matrix $\boldsymbol{D}$ is diagonal with the eigenvalues λ of matrix $\boldsymbol{C}$ on the diagonal:

$$\boldsymbol{D} = \begin{pmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{pmatrix} \qquad \boldsymbol{D}^{-1} = \begin{pmatrix} 1/\lambda_1 & & & \\ & 1/\lambda_2 & & \\ & & \ddots & \\ & & & 1/\lambda_n \end{pmatrix}$$

$$\boxed{\boldsymbol{C}^{-1}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{D}^{-1}} \times \boxed{\boldsymbol{U}^T}$$

$$\boldsymbol{a} = \boldsymbol{U}\boldsymbol{D}^{-1}\boldsymbol{U}^T\boldsymbol{c}\,.$$

Case of singular matrix $\boldsymbol{C}$

At least one of the eigenvalues is zero and the solution above is impossible.

The columns of matrix $\boldsymbol{U}$ (these are the eigenvectors of matrix $\boldsymbol{C}$), which correspond to the *zero* eigenvalues, form an orthogonal basis of the nullspace of $\boldsymbol{C}$. Any arbitrary vector of the null space can be added to the solution $\boldsymbol{a}$ in $\boldsymbol{Ca} = \boldsymbol{c}$.

How to avoid this kind of arbitrariness?

Better define a solution with *no* contribution from the null space: define a modified matrix $\boldsymbol{D}^{*-1}$, where the diagonal elements for zero eigenvalues of $\boldsymbol{D}$ are set to zero (thus replacing $1/0 = \infty$ by 0). The modified solution

$$\boldsymbol{a} = \boldsymbol{U}\boldsymbol{D}^{*-1}\boldsymbol{U}^T\boldsymbol{c}$$

contains *no* contribution from the null space.

It can be shown, that this is the solution with smallest length of $|\boldsymbol{a}|$.

Case of almost singular matrix $\boldsymbol{C}$

Sometimes a system of equation is not singular, but ill conditioned: some of the eigenvalues are very small and the inversion is inaccurate.

This property is visible in the condition number of the matrix, the ratio of the largest to the smallest eigenvalue: $\kappa = \lambda_{\max}/\lambda_{\min}$.

In such a case a solution with modified matrix Matrix $\boldsymbol{D}^{*-1}$ is preferred, where the small eigenvalues are replaced by zero.

Singular value decomposition

A $m \times n$ Matrix $\boldsymbol{A}$ can be expressed in the form

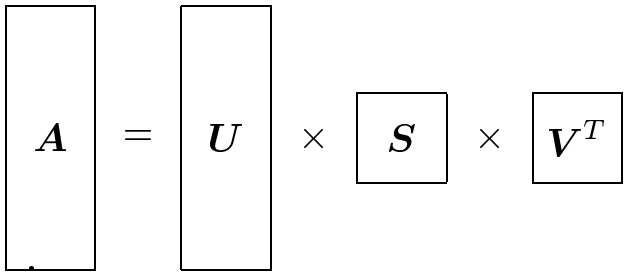

$\boldsymbol{U}$ is a column-orthogonal $m \times n$ matrix: $\boldsymbol{U}^T\boldsymbol{U} = \mathbf{1}$
$\boldsymbol{V}$ is an orthogonal $n \times n$ matrix: $\boldsymbol{V}^T\boldsymbol{V} = \mathbf{1} \quad \boldsymbol{V}\boldsymbol{V}^T = \mathbf{1}$

Matrix S is diagonal and has positive or zero elements, called *singular values.*

Modified solution

Matrix S is diagonal and has positive or zero elements, called *singular values*:

$$\boldsymbol{S} = \begin{pmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{pmatrix} \qquad \boldsymbol{S}^{-1} = \begin{pmatrix} 1/\sigma_1 & & & \\ & 1/\sigma_2 & & \\ & & \ddots & \\ & & & 1/\sigma_n \end{pmatrix}$$

Comparison with previous formalism shows:

$$\boldsymbol{C} = \boldsymbol{A}^T\boldsymbol{A} = \boldsymbol{V}\boldsymbol{S}\boldsymbol{S}\boldsymbol{V}^T \qquad \boldsymbol{D} = \boldsymbol{S}\boldsymbol{S}$$

i.e. eigenvalues of $\boldsymbol{D}$ are $\lambda_i = \sigma_i^2$.
The solution vector is the least squares solution with minimum length

$$\boldsymbol{a} = \boldsymbol{V}\boldsymbol{S}^{*-1}\left(\boldsymbol{U}^T\boldsymbol{y}\right) \ .$$

As above a modified matrix $\boldsymbol{S}^{*-1}$ is defined, where the diagonal elements for small singular values S_{jj} are set to zero.

Comparison

The two methods are equivalent, but there are differences:

- The singular value method is numerically more accurate ($\sigma_i = \sqrt{\lambda_i}$).
- The eigenvalue method works on more compact matrices (small $n \times n$ square matrices instead of large rectangular $m \times n$ matrices).

Both methods avoid those components in the solution vector $\boldsymbol{x}$ from the Null-space, which have no effect on the right hand-side of the initial equations.

Linear least squares – statistical considerations

The model of Linear Least Squares: $\boldsymbol{y} = \boldsymbol{A}\,\boldsymbol{a}$

$$\begin{aligned} \boldsymbol{y} &= \text{vector of measured data} \\ \boldsymbol{A} &= \text{matrix (fixed)} \\ \boldsymbol{a} &= \text{vector of parameters} \\ \boldsymbol{r} &= \boldsymbol{y} - \boldsymbol{A}\boldsymbol{a} = \text{vector of residuals} \\ \boldsymbol{V}[\boldsymbol{y}] &= \text{covariance matrix of the data} \\ \boldsymbol{W} &= \boldsymbol{V}[\boldsymbol{y}]^{-1} \text{ weight matrix} \end{aligned}$$

Least Squares Principle: minimize the expression

$$S(\boldsymbol{a}) = \boldsymbol{r}^T \boldsymbol{W} \boldsymbol{r} = (\boldsymbol{y} - \boldsymbol{A}\boldsymbol{a})^T \; W \; (\boldsymbol{y} - \boldsymbol{A}\boldsymbol{a})$$

with respect to $\boldsymbol{a}$.

Least Squares solution

Derivatives of expression $S(\boldsymbol{a})$:

$$\begin{aligned} \frac{\partial S}{\partial \boldsymbol{a}} &= 2\left(-\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{y} + \left(\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}\right)\boldsymbol{a}\right) \\ \frac{1}{2}\frac{\partial^2 S}{\partial \boldsymbol{a}^2} &= \left(\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}\right) \qquad = \text{constant} \end{aligned}$$

Solution ($\partial S/\partial \boldsymbol{a} = 0$) is linear transformation of the data vector $\boldsymbol{y}$:

$$\hat{\boldsymbol{a}} = \left(\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{W}\,\boldsymbol{y} \qquad = \boldsymbol{B}\boldsymbol{y}$$

Covariance matrix of $\boldsymbol{a}$ by "error" propagation ($\boldsymbol{V}[\boldsymbol{y}] = \boldsymbol{W}^{-1}$):

$$\boldsymbol{V}[\hat{\boldsymbol{a}}] = \boldsymbol{B}\,\boldsymbol{V}[\boldsymbol{y}]\,\boldsymbol{B}^T = \left(\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}\right)^{-1}$$

(identical to inverse of second derivate of S).

Properties of linear least squares solution

Starting from **Principles** properties of the solution are derived, which are valid under certain conditions.

Conditions:

- Data are unbiased: $E[\boldsymbol{y}] = \boldsymbol{A}\,\bar{\boldsymbol{a}}$ ($\bar{\boldsymbol{a}}$ = true parameter vector)
- Covariance matrix $\boldsymbol{V}[\boldsymbol{y}]$ is known (correct) and finite

Properties:

- Estimated parameters are unbiased: $E[\hat{\boldsymbol{a}}] = \bar{\boldsymbol{a}}$

$$E[\hat{\boldsymbol{a}}] = \left(\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{W}\,E[\boldsymbol{y}] = \bar{\boldsymbol{a}}$$

- In the class of unbiased estimates $\boldsymbol{a}^*$, which are linear in the data, the Least Squares estimates $\hat{\boldsymbol{a}}$ have the smallest variance (Gauß-Markoff theorem): $\boldsymbol{V}\left[\boldsymbol{a}^*\right]_{jj} \geq \boldsymbol{V}\left[\hat{\boldsymbol{a}}\right]_{jj}$ for all j
- Expectation of the sum of squares is $\hat{S} = \sigma^2(n-p)$

Properties are not valid, if conditions violated.

Normal equations for uncorrelated data

The diagonal elements of the weight matrix $\boldsymbol{W}$ are $w_i = 1/\sigma_i^2$.

Each data value y_i with its weight w_i makes an independent contribution to the final matrix products.

$$i\text{-th row of } \boldsymbol{A}: \qquad \boldsymbol{A}_i = (d_1, d_2, \ldots, d_p)$$

The contributions of this row to $\boldsymbol{C} = \boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}$ and $\boldsymbol{b} = \boldsymbol{A}^T\boldsymbol{W}\boldsymbol{y}$ can be written as the $p \times p$-matrix $w_i\boldsymbol{A}_i^T \cdot \boldsymbol{A}_i$ and the p-vector $w_i\boldsymbol{A}_i^T \cdot y_i$:

	d_1	d_2	$\ldots$	d_p
d_1	$w_i d_1^2$	$w_i d_1 d_2$	$\ldots$	$w_i d_1 d_p$
d_2		$w_i d_2^2$	$\ldots$	$w_i d_2 d_p$
$\ldots$			$\ldots$	$\ldots$
d_p				$w_i d_p^2$

	y_i
d_1	$w_i d_1 y_i$
d_2	$w_i d_2 y_i$
$\ldots$	$\ldots$
d_p	$w_i d_p y_i$

,

(symmetric elements in the lower half not shown)

Straight line fit

Example: track fit of y (measured) vs. abscissa x

$$y_i = a_0 + a_1 \cdot x_i$$

Matrix $\boldsymbol{A}$ and parameter vector $\boldsymbol{a}$

$$\boldsymbol{A} = \begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix} \qquad \boldsymbol{a} = \begin{pmatrix} a_0 \\ a_1 \end{pmatrix}$$

Weight matrix for *independent* measurements is diagonal (weights $\rightarrow w_i$).

$$\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A} = \begin{pmatrix} \sum w_i & \sum w_i x_i \\ \sum w_i x_i & \sum w_i x_i^2 \end{pmatrix} \qquad \boldsymbol{A}^T\boldsymbol{W}\boldsymbol{y} = \begin{pmatrix} \sum w_i y_i \\ \sum w_i x_i y_i \end{pmatrix}$$

If one measured y_i-value is shifted (biased), then

- parameters biased, and
- χ^2-value high.

Straight line fit

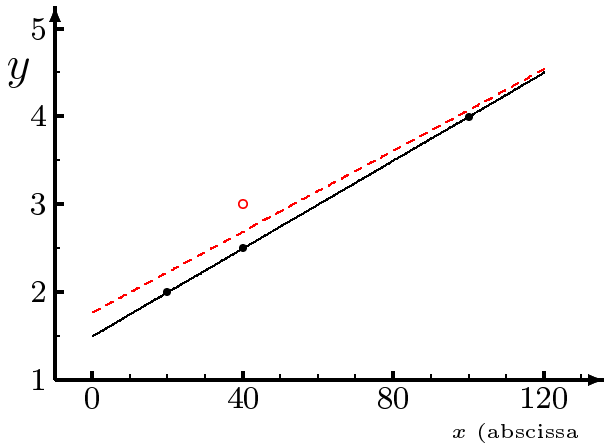

The full line is a straight line fit to three well aligned data points (black dots).

The dashed curve is the straight line fit, if the middle point is "badly aligned" (circle).

Inversion by partitioning

Matrix inversion requires in general $\propto n^3$ operations! Can this be reduced?

A square $n \times n$ matrix $\boldsymbol{A}$ and its inverse $\boldsymbol{B} = \boldsymbol{A}^{-1}$ can be partitioned:

$$\boldsymbol{A} = \left(\begin{array}{c|c} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \hline \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{array}\right) \qquad \boldsymbol{B} = \left(\begin{array}{c|c} \boldsymbol{B}_{11} & \boldsymbol{B}_{12} \\ \hline \boldsymbol{B}_{21} & \boldsymbol{B}_{22} \end{array}\right)$$

$\boldsymbol{A}_{11}$ is a $p \times p$ and $\boldsymbol{A}_{22}$ is a $q \times q$ square matrix $(p + q = n)$.

Submatrices of inverse matrix:

$$\begin{aligned} \boldsymbol{B}_{22} &= \left(\boldsymbol{A}_{22} - \boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\right)^{-1} \\ \boldsymbol{B}_{21} &= -\boldsymbol{B}_{22}\boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1} \qquad\qquad \boldsymbol{B}_{12} = -\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{22} \\ \boldsymbol{B}_{11} &= \boldsymbol{A}_{11}^{-1} - \boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{21} \end{aligned}$$

Large reduction in number of operations if $\boldsymbol{A}_{11}^{-1}$ already known or simple (e.g. diagonal)! Indices 1 and 2 may be interchanged.

Example for partitioned matrices

$$\boldsymbol{C} \cdot \boldsymbol{a} = \boldsymbol{b}$$

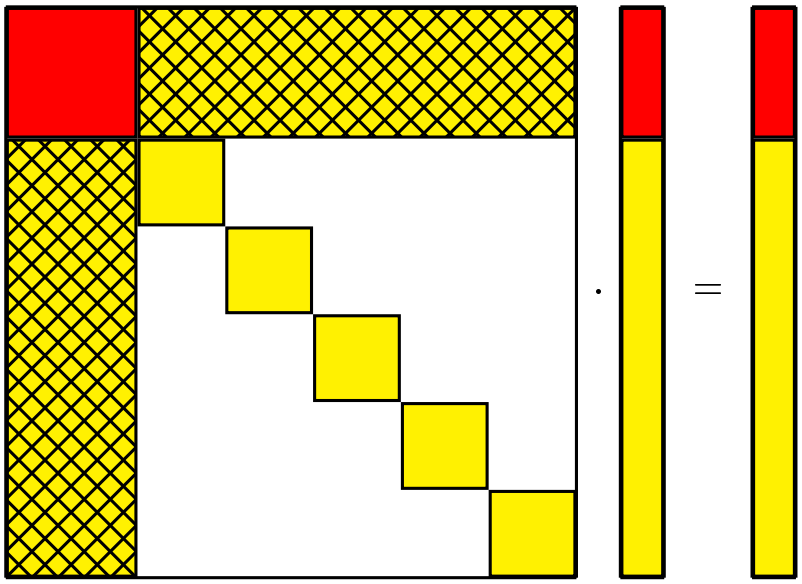

Robust least square fit – outlier treatment

Least square fit on observations y_i, yielding fitted values $\hat{y}_i$, and residuals $r_i = y_i - \hat{y}_i$.

- "Clean" the data by pulling outliers towards their fitted values: winsorize the observations y_i and replace them by pseudo-observations y_i^*:

$$\begin{aligned} y_i^* &= y_i\,, && \text{if} \quad |r_i| \le c\,s_i\,, \\ &= \hat{y}_i - c\,s_i\,, && \text{if} \quad r_i < -c\,s_i\,, \\ &= \hat{y}_i + c\,s_i\,, && \text{if} \quad r_i > +c\,s_i\,. \end{aligned}$$

  (s_i is estimate for the standard error of y_i)
  The factor c regulates the amount of robustness, a goid choice is $c = 1.5$ (note: outliers are not lost!).
- Refit iteratively: the pseudo-observations y_i^* are used to calculate new parameters and new fitted values $\hat{y}_i$.

Example: Outlier treatment in alignment procedure

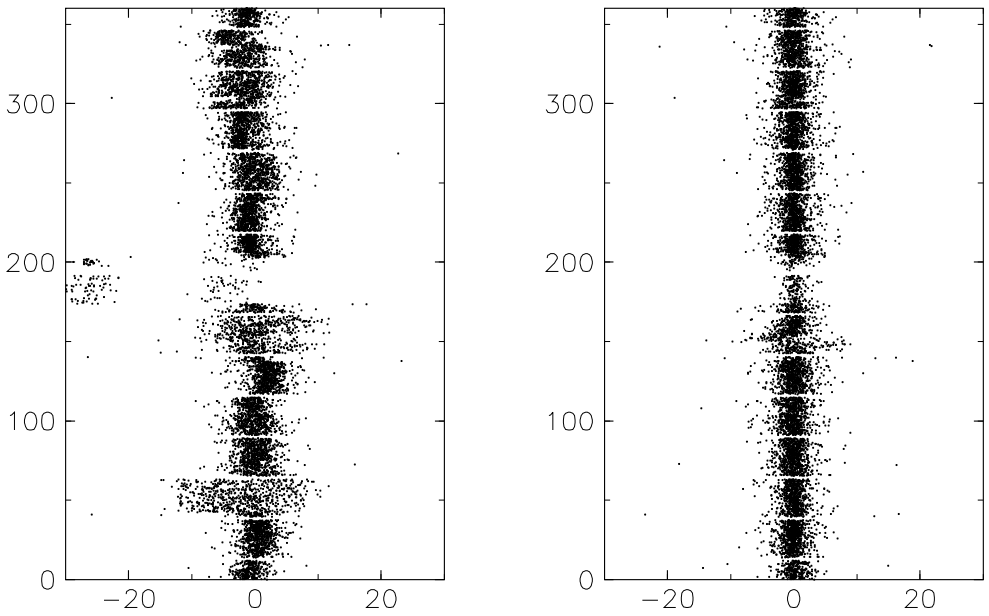

Shown are residuals of a Si tracker as a function of φ in degrees *before* (left) and *after* (right) alignment.

Orthogonal polynomials

If data interpolation required, but no parametrization known: fit of normal general p-th order polynomial to the data

$$f(x_i) \cong y_i \qquad \text{with} \quad f(x) = \sum_{j=0}^{p} a_j \, x^j$$

Generalization of straight-line fit to general p-th order polynomial straightforward: matrix $\boldsymbol{A}^T \boldsymbol{W} \boldsymbol{A}$ contains sum of all powers of x_i up to x_i^{2p}.

Disadvantage:

- fit numerically unstable for larger p
- determination of optimal order p difficult
- value of coefficients a_j depend on highest exponent p

Better: fit of orthogonal polynomial.

Straight line fit

Orthogonal parametrization

$$f(x) = a_0 + a_1 x \quad \text{replaced by} \quad f(x) = a_0 \cdot p_0(x) + a_1 \cdot p_1(x)$$

$$\begin{aligned} p_0(x) &= b_{00} \\ p_1(x) &= b_{10} + b_{11} x = b_{11}\left(x - \langle x \rangle\right) \quad . \end{aligned}$$

$$\text{with} \qquad b_{00} = \left(\sum_{i=1}^{n} w_i \right)^{-1/2} \qquad b_{11} = \left(\sum_{i=1}^{n} w_i \left(x_i - \langle x \rangle\right)^2 \right)^{-1/2}$$

and with the coefficient $b_{10} = -\langle x \rangle\, b_{11}$.

$$\begin{aligned} \boldsymbol{A}^T \boldsymbol{W} \boldsymbol{A} &= \begin{pmatrix} \sum w_i p_0^2(x_i) & \sum w_i p_0(x_i) p_1(x_i) \\ \sum w_i p_0(x_i) p_1(x_i) & \sum w_i p_1^2(x_i) \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ \boldsymbol{A}^T \boldsymbol{W} \boldsymbol{y} &= \begin{pmatrix} \sum w_i p_0(x_i) y_i \\ \sum w_i p_1(x_i) y_i \end{pmatrix} = \begin{pmatrix} a_0 \\ a_1 \end{pmatrix} \end{aligned}$$

i.e. only sums, no matrix inversion required.

Construction of higher order polynomial $p_j(x)$

General orthogonal polynomial: $$f(x) = \sum_{j=1}^{p} a_j p_j(x)$$

$p_j(x)$ constructed from the two previous functions $p_{j-1}(x)$ and $p_{j-2}(x)$ with parameters α, β and γ given by sums:

Recurrence relation $$\boxed{\gamma p_j(x) = (x_i - \alpha)\, p_{j-1}(x) - \beta p_{j-2}(x)}$$

Parameters $\widehat{a}_j$ are determined from data by

$$\widehat{a}_j = \sum_{i=1}^{n} w_i p_j(x_i)\; y_i$$

Covariance matrix is **1**, i.e. all variances are 1 and the coefficients are *independent*.

Example: Data points and third order orthogonal polynomia

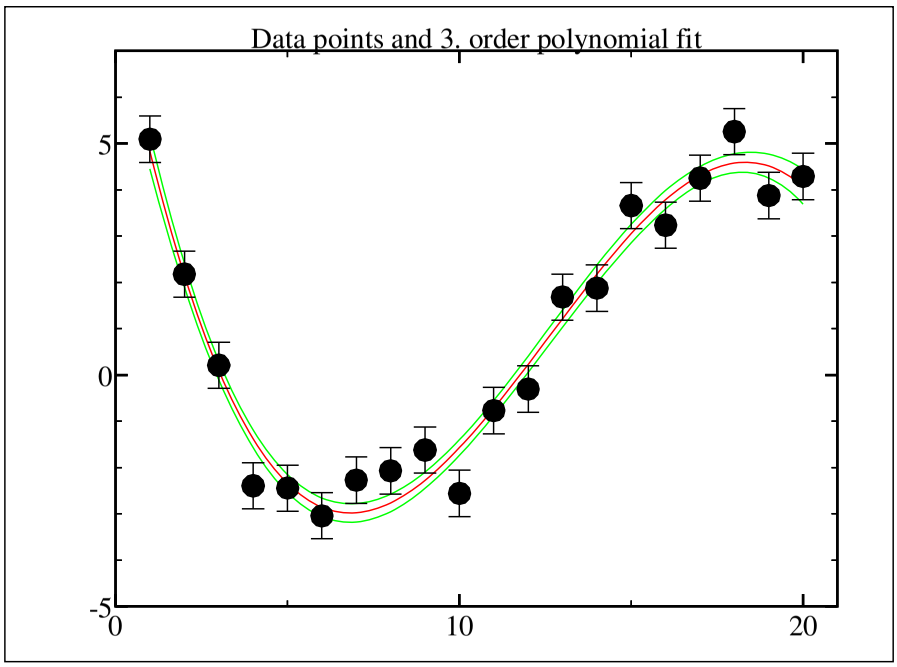

Shown is polynomial (red) and ± 1 standard deviation curves (green).

Example: Coefficients $a_0 \ldots a_{10}$

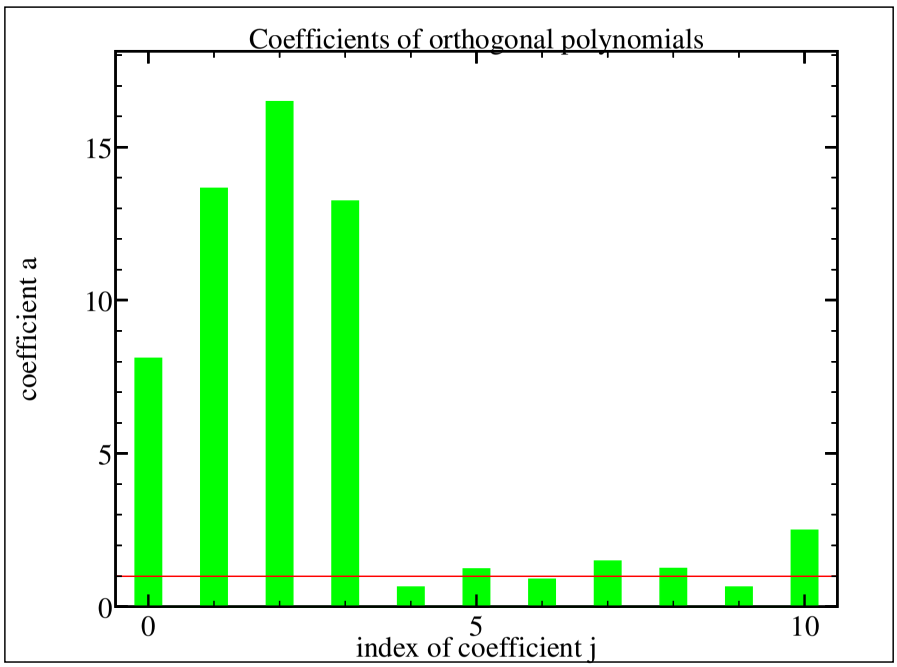

Coefficients $a_4 \ldots a_{10}$ compatible with $E[a_j] = 0$ (noise).

Least squares with constraints

Aim: improve measured values using known relations between data.

Linear or non-linear (equality) constraints:

$$\boxed{f_k(\boldsymbol{a}_{\text{true}}, \boldsymbol{y}_{\text{true}}) = 0 \qquad k = 1,\, 2, \ldots m}$$

with

$$\begin{aligned} \boldsymbol{y} &= n\text{-vector of measured data} \\ \boldsymbol{V}[\boldsymbol{y}] &= \text{covariance matrix of the data} \\ \boldsymbol{W} &= \boldsymbol{V}[\boldsymbol{y}]^{-1} \text{ weight matrix} \\ \boldsymbol{a} &= p\text{-vector of parameters, unmeasured} \end{aligned}$$

Least squares principle: minimize with respect to $\Delta\boldsymbol{y}$ and $\boldsymbol{a}$

$$S(\boldsymbol{y}) = \Delta\boldsymbol{y}^T \boldsymbol{W} \Delta\boldsymbol{y}$$

under the conditions

$$f_k(\boldsymbol{a}, \boldsymbol{y} + \Delta\boldsymbol{y}) = 0 \quad k = 1,\, 2, \ldots m$$

Kinematical constraints

Example: reconstruction of kinematic variables is often overconstrained – kinematical fit possible.

A different analysis style is standard at DESY:

- measurement errors are ignored,
- several formulas are used, based on subsets of data (no fit),
- comparison of alternatives (electron method, hadron method, double angle method, Σ method) by MC.

e.g. $\quad y_{\text{el}} = 1 - \dfrac{E'_e}{2E_e}\,(1 - \cos\theta'_e) \qquad y_{\text{JB}} = \dfrac{\delta_h}{2E_e}$

$$y_{\text{DA}} = \frac{\sin\theta'_e\,(1 - \cos\gamma_h)}{\sin\theta'_e + \sin\gamma_h - \sin(\theta'_e + \gamma_h)} \qquad y_\Sigma = \frac{\delta_h}{\delta}$$

Should there be an optimal method, which uses all measured data?

Method of Lagrange multipliers

Additional m parameters λ_k, the Lagrange multipliers, are introduced, one for each equation and a new function is defined:

$$L(\boldsymbol{a}, \Delta\boldsymbol{y}) = S(\Delta\boldsymbol{y}) + 2\sum_{k=1}^{m} \lambda_k f_k(\boldsymbol{a}, \boldsymbol{y} + \Delta\boldsymbol{y})$$

The necessary condition for a local extremum of this function with respect to all parameters $\Delta\boldsymbol{y}$, $\boldsymbol{a}$ and λ is equivalent to the condition of a minimum of $S(\boldsymbol{a}, \Delta\boldsymbol{y})$ under conditions $f_k(\boldsymbol{a}, \boldsymbol{y} + \Delta\boldsymbol{y}) = 0$.

Number of unknowns

n	measured parameters and corrections Δy_i
p	parameters a_j
m	multipliers λ_k
$n+p+m$	parameters in total

Linearization – vector notation

$$\boldsymbol{f} + \boldsymbol{A}(\Delta\boldsymbol{a} - \Delta\boldsymbol{a}^*) + \boldsymbol{B}(\Delta\boldsymbol{y} - \Delta\boldsymbol{y}^*) = 0$$

or $\boxed{\boldsymbol{A}\Delta\boldsymbol{a} + \boldsymbol{B}\Delta\boldsymbol{y} = \boldsymbol{c}}$ with $\boldsymbol{c} = \boldsymbol{A}\Delta\boldsymbol{a}^* + \boldsymbol{B}\Delta\boldsymbol{y}^* - \boldsymbol{f}$

$$\boldsymbol{A} = \begin{pmatrix} \partial f_1/\partial a_1 & \partial f_1/\partial a_2 & \dots & \partial f_1/\partial a_p \\ \partial f_2/\partial a_1 & \partial f_2/\partial a_2 & \dots & \partial f_2/\partial a_p \\ & & & \\ \dots & & & \\ \partial f_m/\partial a_1 & \partial f_m/\partial a_2 & \dots & \partial f_m/\partial a_p \end{pmatrix} \qquad \boldsymbol{f} = \begin{pmatrix} f_1(\boldsymbol{a}^*, \boldsymbol{y}^*) \\ f_2(\boldsymbol{a}^*, \boldsymbol{y}^*) \\ \\ \dots \\ f_m(\boldsymbol{a}^*, \boldsymbol{y}^*) \end{pmatrix}$$

$$\boldsymbol{B} = \begin{pmatrix} \partial f_1/\partial y_1 & \partial f_1/\partial y_2 & \dots & \partial f_1/\partial y_n \\ \partial f_2/\partial y_1 & \partial f_2/\partial y_2 & \dots & \partial f_2/\partial y_n \\ \dots & & & \\ \partial f_m/\partial y_1 & \partial f_m/\partial y_2 & \dots & \partial f_m/\partial y_n \end{pmatrix}$$

($\Delta\boldsymbol{y}^*$ and $\Delta\boldsymbol{a}^*$ are previous corrections, initially 0)

Linearization

Function L with linearization of constraints:

$$L = \Delta \boldsymbol{y}^T \boldsymbol{W} \Delta \boldsymbol{y} + 2\lambda^T (\boldsymbol{A}\Delta \boldsymbol{a} + \boldsymbol{B}\Delta \boldsymbol{y} - \boldsymbol{c})$$

The necessary conditions for an extremum are obtained by differentiation:

$$\begin{array}{lllll} \boldsymbol{W}\Delta \boldsymbol{y} & & + \boldsymbol{B}^T\lambda & = & 0 \\ & & \boldsymbol{A}^T\lambda & = & 0 \\ \boldsymbol{B}\Delta \boldsymbol{y} + \boldsymbol{A}\Delta \boldsymbol{a} & & & = & \boldsymbol{c} \end{array}$$

This system of coupled matrix equations (in total $n + p + m$ equations) has to be solved for the unknowns $\Delta \boldsymbol{y}$, $\Delta \boldsymbol{a}$ and λ.

Case of no unknowns

$$\begin{array}{lllll} \boldsymbol{W}\Delta \boldsymbol{y} & + \boldsymbol{B}^T\lambda & = & 0 \\ \boldsymbol{B}\Delta \boldsymbol{y} & & = & \boldsymbol{c} \end{array}$$

Solution: multiply first equation with $\boldsymbol{W}^{-1}$

$$\ldots \text{from the left} \qquad \Delta \boldsymbol{y} = -\boldsymbol{W}^{-1}\boldsymbol{B}^T\lambda \qquad (*)$$

$$\text{insert into second equation} \qquad \left(\boldsymbol{B}\boldsymbol{W}^{-1}\boldsymbol{B}^T\right)\lambda = -\boldsymbol{c}$$

$$\text{solve for } \boldsymbol{\lambda} \text{ by} \qquad \lambda = -\boldsymbol{W_B}\,\boldsymbol{c} \quad \text{with} \quad \boldsymbol{W}_B = \left(\boldsymbol{B}\boldsymbol{W}^{-1}\boldsymbol{B}^T\right)^{-1}$$

$$\text{insert into equation } (*) \qquad \Delta \boldsymbol{y} = \left(\boldsymbol{W}^{-1}\boldsymbol{B}^T\boldsymbol{W}_B\right)\,\boldsymbol{c}$$

For *linear* problems with constant $\boldsymbol{B}$ this is the solution; for *non-linear* problems iterations are necessary.

$$\hat{\boldsymbol{y}} = \boldsymbol{y} + \Delta \boldsymbol{y} \qquad \boldsymbol{V}(\hat{\boldsymbol{y}}) = \boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\left(\boldsymbol{B}^T\boldsymbol{W}_B\boldsymbol{B}\right)\boldsymbol{W}^{-1}$$

with covariance matrix by error propagation.

Pulls = normalized corrections

On average all corrections Δy_i have the same magnitude w.r.t. their standard deviation. A check of systematic deviations is possible by looking at distributions of pulls (= normalized corrections).

Covariance matrix of the corrections $\Delta \boldsymbol{y}$:

$$\boldsymbol{V}(\Delta \boldsymbol{y}) = \boldsymbol{W}^{-1}\left(\boldsymbol{B}^T \boldsymbol{W}_B \boldsymbol{B}\right)\boldsymbol{W}^{-1} = \boldsymbol{V}(\boldsymbol{y}) - \boldsymbol{V}(\hat{\boldsymbol{y}})$$

Normalized *pull values* are obtained by

$$p_i = \frac{\Delta y_i}{\sqrt{V(\boldsymbol{y})_{ii} - V(\hat{\boldsymbol{y}})_{ii}}}$$

The pulls should follow the standardized Gaussian distribution $N(0,1)$, if the measured data are normally distributed and the conditions are linear.

The general case

System of equations to be solved in the general case:

$$\begin{pmatrix} \boldsymbol{W} & 0 & \boldsymbol{B}^T \\ 0 & 0 & \boldsymbol{A}^T \\ \boldsymbol{B} & \boldsymbol{A} & 0 \end{pmatrix} \begin{pmatrix} \Delta \boldsymbol{y} \\ \Delta \boldsymbol{a} \\ \lambda \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \boldsymbol{c} \end{pmatrix}$$

What is the difference between *measured* and *unmeasured* parameters?

Elements of weight matrix w.r.t. to unmeasured parameters is zero!

Inverse of partitioned matrix:

$$\begin{pmatrix} \boldsymbol{W} & 0 & \boldsymbol{B}^T \\ 0 & 0 & \boldsymbol{A}^T \\ \boldsymbol{B} & \boldsymbol{A} & 0 \end{pmatrix}^{-1} = \begin{pmatrix} \boldsymbol{C}_{11} & \boldsymbol{C}_{21}^T & \boldsymbol{C}_{31}^T \\ \boldsymbol{C}_{21} & \boldsymbol{C}_{22} & \boldsymbol{C}_{32}^T \\ \boldsymbol{C}_{31} & \boldsymbol{C}_{32} & \boldsymbol{C}_{33} \end{pmatrix}$$

Solution

Elements of the inverse matrix:

$$\begin{aligned}
\boldsymbol{C}_{11} &= \boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\boldsymbol{B}^T\boldsymbol{W}_B\boldsymbol{B}\boldsymbol{W}^{-1} \\
&\quad +\boldsymbol{W}^{-1}\boldsymbol{B}^T\boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B\boldsymbol{B}\boldsymbol{W}^{-1} \\
\boldsymbol{C}_{21} &= -\boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B\boldsymbol{B}\boldsymbol{W}^{-1} \\
\boldsymbol{C}_{22} &= \boldsymbol{W}_A^{-1} \\
\boldsymbol{C}_{31} &= \boldsymbol{W}_B\boldsymbol{B}\boldsymbol{W}^{-1} - \boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B\boldsymbol{B}\boldsymbol{W}^{-1} \\
\boldsymbol{C}_{32} &= \boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1} \\
\boldsymbol{C}_{33} &= -\boldsymbol{W}_B + \boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B
\end{aligned}$$

with abbreviations

$$\begin{aligned}
\boldsymbol{W}_B &= (\boldsymbol{B}\boldsymbol{W}^{-1}\boldsymbol{B}^T)^{-1} \\
\boldsymbol{W}_A^{-1} &= (\boldsymbol{A}^T\boldsymbol{W}_B\boldsymbol{A})^{-1}
\end{aligned}$$

Corrections, covariance matrix and checks

Corrections $\Delta\boldsymbol{y}$, $\Delta\boldsymbol{a}$ and the Lagrangian multipliers λ are obtained by multiplication:

$$\begin{aligned}
\Delta\boldsymbol{y} &= \boldsymbol{C}_{31}^T\,\boldsymbol{c} &&= (\boldsymbol{W}^{-1}\boldsymbol{B}^T\boldsymbol{W}_B - \boldsymbol{W}^{-1}\boldsymbol{B}^T\boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1}\boldsymbol{A}\boldsymbol{W}_B)\,\boldsymbol{c} \\
\Delta\boldsymbol{a} &= \boldsymbol{C}_{32}^T\,\boldsymbol{c} &&= \boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B\,\boldsymbol{c} \\
\lambda &= \boldsymbol{C}_{33}\,\boldsymbol{c} &&= (-\boldsymbol{W}_B + \boldsymbol{W}_B\boldsymbol{A}\boldsymbol{W}_A^{-1}\boldsymbol{A}^T\boldsymbol{W}_B)\,\boldsymbol{c}
\end{aligned}$$

Covariance matrix of the combined vector $\hat{\boldsymbol{y}}$, $\hat{\boldsymbol{a}}$ is

$$\boldsymbol{V}\begin{pmatrix}\hat{\boldsymbol{y}}\\ \hat{\boldsymbol{a}}\end{pmatrix} = \begin{pmatrix}\boldsymbol{C}_{11} & \boldsymbol{C}_{21}^T\\ \boldsymbol{C}_{21} & \boldsymbol{C}_{22}\end{pmatrix}$$

Checks: $E(\hat{S})$ and pulls:
Value of squares of deviations: $E(\hat{S}) = (m-p)$, i.e. on average 1 per degree of freedom, with

$$\hat{S} \quad \text{distributed} \quad \chi^2_{m-p}$$

if the measured data are Gaussian distributed and the conditions are linear.

Code for $\pi^0 \to \gamma\gamma$ fit using APLCON

```
      VX = ...     ! joint covariance matrix
*     6 parameters and 1 constraint equations
      CALL SIMNIT(6,1,IDEB,1.E-5)
*     add four vectors
   20 EN = X(1) + X(4)
      PX = X(1)*SIN(X(2))*COS(X(3)) + X(4)*SIN(X(5))*COS(X(6))
      PY = X(1)*SIN(X(2))*SIN(X(3)) + X(4)*SIN(X(5))*SIN(X(6))
      PZ = X(1)*COS(X(2))           + X(4)*COS(X(5))
*     mass squared by square of four vector
      FM2  = EN**2 - PX**2 - PY**2 - PZ**2
*     constraint = calculated mass**2 - pi zero mass**2
      F(1) = FM2 - 0.135**2
*     call fit program
      CALL APLCON(X,VX,F,IREP)
      IF(IREP.LT.0) GOTO 20
      PMF=SQRT(PX**2+PY**2+PZ**2)
      XMF=SQRT(FM2)
```


Alignment Problems

Example: z-coordinate measured in detector with many (e.g. ten) parallel planes, with model

$$z_i = z_0 + (\cot\vartheta) \cdot S_i$$

Some or all planes may be misaligned (Δz), visible in bad χ^2 of fits.

Standard method in HEP based on residuals:

- make straight line fits,
- plot histograms of residuals for all planes,
- determine in all histograms mean residuum $\Rightarrow \Delta z$ correction

With corrected z-values: the χ^2 of fits is much smaller!

$\Rightarrow$ alignment called "successful"!

Is this alignment really successful?

Although the χ^2 of fits is much smaller, the parameter values are almost unchanged and still biased!

...the original fit result was biased, because a wrong model (without shifts) has been used.

Further properties of method based on residuals:

Corrections Δz_j

- depend on chamber (in)efficiencies,
- are unstable, if "alignment" is done repeatedly.

Improved Alignment

Improvement I: try to use correct model

$$z_i = z_0 + (\cot\vartheta) \cdot S_i + \Delta z_j$$

where Δz_j = z-correction of plane j.

Problem: 10 measured values z_i, but $2 + 10$ parameters to be fitted (no solution).

Improvement II: use correct model simultaneously for N tracks.

$10 \times N$ measured values, and $2 \times N + 10$ parameters $\Rightarrow$ fit possible, for example $N = 1000$: 10 000 measured values 2010 parameters

This model is still unstable – matrix $\left(\boldsymbol{A}^T \boldsymbol{W} \boldsymbol{A}\right)$ is singular with rank defect 2: overall shift plus rotation of detector not fixed!

Improvement III: add two constraints for shifts Δz_j (simple in simultaneous fit, difficult in residual method).

MILLEPEDE – *a Fit Program for many + x Parameters*

MILLEPEDE – a general program for Linear Least Squares (LLS) with a large number of parameters, especially for alignment and calibration problems.

LLS with solution $\boldsymbol{a} = \boldsymbol{C}^{-1}\,\boldsymbol{A}^T\boldsymbol{W}\boldsymbol{y}$ requires inversion of symmetric matrix

$$\boldsymbol{C} = \boldsymbol{A}^T\boldsymbol{W}\boldsymbol{A}$$

Execution time $\propto n^3$ for $n \times n$ matrix, memory $\propto n^2$.

$\Rightarrow$ Practical limit around $n = 1000$: time $\approx$ few minutes, 4 Mbytes for matrix $\boldsymbol{C}$ (double precision).

Matrix often singular or close to singular $\Rightarrow$ special inversion routine for symmetric matrices, which determines invertible part of the matrix (remaining elements and corresponding parameters $\Rightarrow$ zero).

Code and manual on homepage: `http://www.desy.de/~blobel/`

Alignment/Calibration Problems

Example: tracks in z-chambers with $z_i = z_0 + (\cot\vartheta) \cdot S_i + \Delta z_j$

Often large number (≈ 1000) of global parameters

plus

large number N of tracks ($\approx 10\,000$) with e.g. $2 \times N$ local parameters.

$\Rightarrow$ standard solution not possible, but:

Interest only in global parameters: huge matrix C can be separated into submatrices related to

- global parameters and
- local parameters,

and due to special structure, can be reduced to system of equations for global parameters only!

System of normal equations

The simultaneous fit of global and local parameters:

$$\left(\begin{array}{c|c|c|c} \sum \boldsymbol{C}_i & \cdots & \boldsymbol{G}_i & \cdots \\ \hline \vdots & \ddots & 0 & 0 \\ \hline \boldsymbol{G}_i^T & 0 & \boldsymbol{\Gamma}_i & 0 \\ \hline \vdots & 0 & 0 & \ddots \end{array}\right) \cdot \left(\begin{array}{c} \boldsymbol{a} \\ \hline \vdots \\ \hline \boldsymbol{\alpha}_i \\ \hline \vdots \end{array}\right) = \cdot \left(\begin{array}{c} \sum \boldsymbol{b}_i \\ \hline \vdots \\ \hline \boldsymbol{\beta}_i \\ \hline \vdots \end{array}\right)$$

Reduction of the Matrix Size

$$\boldsymbol{C} \cdot \boldsymbol{a} = \boldsymbol{b}$$

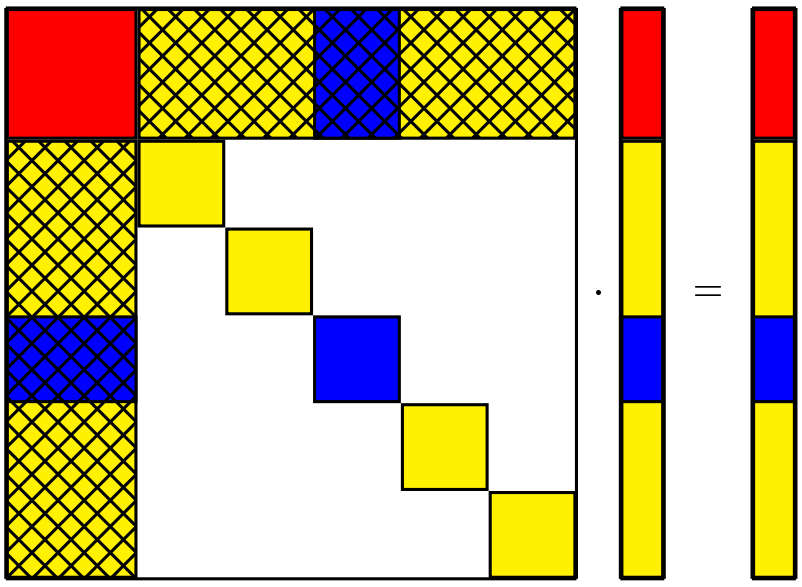

red = global parameters yellow (and blue) = local parameters

Reduction of matrix size

The matrices $\boldsymbol{C}_i$, $\boldsymbol{\Gamma}_i$, $\boldsymbol{G}_i$ and the vectors $\boldsymbol{b}_i$ and $\boldsymbol{\beta}_i$ contain contributions from the i-th track. The only connection between the local parameters of different tracks is given via the sub-matrices $\boldsymbol{G}_i$ und $\boldsymbol{C}_i$.

Ignoring the global parameters one could solve the normal equations $\boldsymbol{\Gamma}_i\boldsymbol{\alpha}_i = \boldsymbol{\beta}_i$ for each track separately by $\boldsymbol{\alpha}_i = \boldsymbol{\Gamma}_i^{-1}\boldsymbol{\beta}_i$.

Huge matrix is reduced to n normal equations for global parameters

$$\begin{pmatrix} & \boldsymbol{C'} & \end{pmatrix} \begin{pmatrix} \\ \boldsymbol{a} \\ \\ \end{pmatrix} = \begin{pmatrix} \\ \boldsymbol{b'} \\ \\ \end{pmatrix} ,$$

with a modified matrix $\boldsymbol{C'}$ and a modified vector $\boldsymbol{b'}$,

$$\boldsymbol{C'} = \sum_i \boldsymbol{C}_i - \sum_i \boldsymbol{G}_i\boldsymbol{\Gamma}_i^{-1}\boldsymbol{G}_i^T \qquad \boldsymbol{b'} = \sum_i \boldsymbol{b}_i - \sum_i \boldsymbol{G}_i \left(\boldsymbol{\Gamma}_i^{-1}\boldsymbol{\beta}_i\right) .$$

System of normal equations

The solution

$$\boldsymbol{a} = \boldsymbol{C'}^{-1}\,\boldsymbol{b'}$$

with covariance matrix $\boldsymbol{C'}^{-1}$ is direct, no iterations are required.

Iterations required however because of outliers (e.g. bad tracks), with loose cuts initially and tight final cuts.

Applications in alignment of trackers:

- H1 Cst Silicon tracker and Cjc Central Jet Chamber
- HeraB
- CDF at Fermilab