

Fitting Linear Statistical Models to Data by Least Squares: Introduction

Radu Balan, Brian R. Hunt and C. David Levermore
University of Maryland, College Park

University of Maryland, College Park, MD

Math 420: *Mathematical Modeling*

January 29, 2020 version

© 2020 R. Balan, B.R. Hunt and C.D. Levermore

Outline

- 1) Introduction to Linear Statistical Models
- 2) Linear Euclidean Least Squares Fitting
- 3) Linear Weighted Least Squares Fitting
- 4) Least Squares Fitting for Univariate Polynomial Models
- 5) Least Squares Fitting with Orthogonalization
- 6) Multivariate Linear Least Squares Fitting
- 7) General Multivariate Linear Least Squares Fitting

1. Introduction to Linear Statistical Models

In modeling one is often faced with the problem of fitting data with some analytic expression. Let us suppose that we are studying a phenomenon that evolves over time. Given a set of n times $\{t_j\}_{j=1}^n$ such that at each time t_j we take a measurement y_j of the phenomenon. We can represent this data as the set of ordered pairs

$$\{(t_j, y_j)\}_{j=1}^n .$$

Each y_j might be a single number or a vector of numbers. For simplicity, we will first treat the *univariate* case when it is a single number. The more complicated *multivariate* case when it is a vector will be treated later.

1. Introduction to Linear Statistical Models

In modeling one is often faced with the problem of fitting data with some analytic expression. Let us suppose that we are studying a phenomenon that evolves over time. Given a set of n times $\{t_j\}_{j=1}^n$ such that at each time t_j we take a measurement y_j of the phenomenon. We can represent this data as the set of ordered pairs

$$\{(t_j, y_j)\}_{j=1}^n.$$

Each y_j might be a single number or a vector of numbers. For simplicity, we will first treat the *univariate* case when it is a single number. The more complicated *multivariate* case when it is a vector will be treated later.

The basic problem we will examine is the following.

How can this data set be used to make a reasonable guess about what a measurement of this phenomenon might yield at any other time?

Model Complexity and Overfitting

Of course, you can always find functions $f(t)$ such that $y_j = f(t_j)$ for every $j = 1, \dots, n$. For example, you can use Lagrange interpolation to construct a unique polynomial of degree at most $n - 1$ that does this. However, such a polynomial often exhibits wild oscillations that make it a useless fit. This phenomena is called *overfitting*.

Model Complexity and Overfitting

Of course, you can always find functions $f(t)$ such that $y_j = f(t_j)$ for every $j = 1, \dots, n$. For example, you can use Lagrange interpolation to construct a unique polynomial of degree at most $n - 1$ that does this. However, such a polynomial often exhibits wild oscillations that make it a useless fit. This phenomena is called *overfitting*.

There are two reasons why such difficulties arise.

- The times t_j and measurements y_j are subject to error, so finding a function that fits the data exactly is not a good strategy.

Model Complexity and Overfitting

Of course, you can always find functions $f(t)$ such that $y_j = f(t_j)$ for every $j = 1, \dots, n$. For example, you can use Lagrange interpolation to construct a unique polynomial of degree at most $n - 1$ that does this. However, such a polynomial often exhibits wild oscillations that make it a useless fit. This phenomena is called *overfitting*.

There are two reasons why such difficulties arise.

- The times t_j and measurements y_j are subject to error, so finding a function that fits the data exactly is not a good strategy.
- The assumed form of $f(t)$ might be ill suited for matching the behavior of the phenomenon over the time interval being considered.

Model Fitting

One strategy to help avoid these difficulties is to draw $f(t)$ from a family of suitable functions, which is called a *model* in statistics. If we denote this model by $f(t; \beta_1, \dots, \beta_m)$ where $m \ll n$ then the idea is to find values of β_1, \dots, β_m such that the graph of $f(t; \beta_1, \dots, \beta_m)$ best fits the data. More precisely, we will define the *residuals* $r_j(\beta_1, \dots, \beta_m)$ by the relation

$$y_j = f(t_j; \beta_1, \dots, \beta_m) + r_j(\beta_1, \dots, \beta_m), \quad \text{for every } j = 1, \dots, n,$$

and try to minimize the $r_j(\beta_1, \dots, \beta_m)$ in some sense.

Model Fitting

One strategy to help avoid these difficulties is to draw $f(t)$ from a family of suitable functions, which is called a *model* in statistics. If we denote this model by $f(t; \beta_1, \dots, \beta_m)$ where $m \ll n$ then the idea is to find values of β_1, \dots, β_m such that the graph of $f(t; \beta_1, \dots, \beta_m)$ best fits the data. More precisely, we will define the *residuals* $r_j(\beta_1, \dots, \beta_m)$ by the relation

$$y_j = f(t_j; \beta_1, \dots, \beta_m) + r_j(\beta_1, \dots, \beta_m), \quad \text{for every } j = 1, \dots, n,$$

and try to minimize the $r_j(\beta_1, \dots, \beta_m)$ in some sense.

The problem is simplified by restricting ourselves to models in which the parameters appear linearly — so-called *linear models*. Such a model is specified by the choice of a basis $\{f_i(t)\}_{i=1}^m$ and takes the form

$$f(t; \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i f_i(t).$$

Polynomial and Periodic Models

Example. The most classic linear model is the family of all *polynomials* of degree less than m . This family is often expressed as

$$f(t; \beta_0, \dots, \beta_{m-1}) = \sum_{i=0}^{m-1} \beta_i t^i.$$

Notice that here the index i runs from 0 to $m - 1$ rather than from 1 to m . This indexing convention is used for polynomial models because it matches the degree of each term in the sum.

Polynomial and Periodic Models

Example. The most classic linear model is the family of all *polynomials* of degree less than m . This family is often expressed as

$$f(t; \beta_0, \dots, \beta_{m-1}) = \sum_{i=0}^{m-1} \beta_i t^i.$$

Notice that here the index i runs from 0 to $m - 1$ rather than from 1 to m . This indexing convention is used for polynomial models because it matches the degree of each term in the sum.

Example. If the underlying phenomena is *periodic* with period T then a classic linear model is the family of all *trigonometric polynomials* of degree at most L . This family can be expressed as

$$f(t; \alpha_0, \dots, \alpha_L, \beta_1, \dots, \beta_L) = \alpha_0 + \sum_{k=1}^L (\alpha_k \cos(k\omega t) + \beta_k \sin(k\omega t)),$$

where $\omega = 2\pi/T$ its *fundamental frequency*. Note that here $m = 2L + 1$.

Translation Invariant Models

Remark. Linear models are linear in the parameters, but are typically nonlinear in the independent variable t . This is illustrated by the foregoing examples: the family of all polynomials of degree less than m is nonlinear in t for $m > 2$; the family of all trigonometric polynomials of degree at most L is nonlinear in t for $L > 0$.

Translation Invariant Models

Remark. Linear models are linear in the parameters, but are typically nonlinear in the independent variable t . This is illustrated by the foregoing examples: the family of all polynomials of degree less than m is nonlinear in t for $m > 2$; the family of all trigonometric polynomials of degree at most L is nonlinear in t for $L > 0$.

Remark. When there is no preferred instant of time it is best to pick a model $f(t; \beta_1, \dots, \beta_m)$ that is *translation invariant*. This means for every choice of parameter values $(\beta_1, \dots, \beta_m)$ and time shift s there exist parameter values $(\beta'_1, \dots, \beta'_m)$ such that

$$f(t + s; \beta_1, \dots, \beta_m) = f(t; \beta'_1, \dots, \beta'_m) \quad \text{for every } t.$$

Both models given on the previous slide are translation invariant. Can you show this? Can you find models that are not translation invariant?

Linear Models

It is as easy to work in the more general setting in which we are given data

$$\{(\mathbf{x}_j, y_j)\}_{j=1}^n,$$

where the \mathbf{x}_j lie within a bounded domain $\mathbb{X} \subset \mathbb{R}^d$ and the y_j lie in \mathbb{R} . We call \mathbf{x} the *independent variable* and y the *dependent variable*.

Linear Models

It is as easy to work in the more general setting in which we are given data

$$\{(\mathbf{x}_j, y_j)\}_{j=1}^n,$$

where the \mathbf{x}_j lie within a bounded domain $\mathbb{X} \subset \mathbb{R}^d$ and the y_j lie in \mathbb{R} . We call \mathbf{x} the *independent variable* and y the *dependent variable*.

The problem we will examine now becomes the following.

How can this data set be used to make a reasonable guess about the value of y when \mathbf{x} is any point in \mathbb{X} ?

Linear Models

It is as easy to work in the more general setting in which we are given data

$$\{(\mathbf{x}_j, y_j)\}_{j=1}^n,$$

where the \mathbf{x}_j lie within a bounded domain $\mathbb{X} \subset \mathbb{R}^d$ and the y_j lie in \mathbb{R} . We call \mathbf{x} the *independent variable* and y the *dependent variable*.

The problem we will examine now becomes the following.

How can this data set be used to make a reasonable guess about the value of y when \mathbf{x} is any point in \mathbb{X} ?

Linear Models

We will consider linear statistical models with m real parameters in the form

$$f(\mathbf{x}; \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i f_i(\mathbf{x}),$$

where each basis function $f_i(\mathbf{x})$ is defined over \mathbb{X} and takes values in \mathbb{R} .

Linear Models

We will consider linear statistical models with m real parameters in the form

$$f(\mathbf{x}; \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i f_i(\mathbf{x}),$$

where each basis function $f_i(\mathbf{x})$ is defined over \mathbb{X} and takes values in \mathbb{R} .

Example. A classic model in this setting is the family of all affine functions. If x_i denotes the i^{th} entry of \mathbf{x} then this family can be written as

$$f(\mathbf{x}; a, b_1, \dots, b_d) = a + \sum_{i=1}^d b_i x_i.$$

Alternatively, it can be expressed in vector notation as

$$f(\mathbf{x}; a, \mathbf{b}) = a + \mathbf{b} \cdot \mathbf{x},$$

where $a \in \mathbb{R}$ and $\mathbf{b} \in \mathbb{R}^d$. In this case $m = d + 1$.

Linear Models

Example. Similarly, the family of all quadratic functions can be expressed in vector notation as

$$f(\mathbf{x}; \mathbf{a}, \mathbf{b}, \mathbf{C}) = a + \mathbf{b} \cdot \mathbf{x} + \mathbf{x} \cdot \mathbf{C}\mathbf{x},$$

where $a \in \mathbb{R}$, $\mathbf{b} \in \mathbb{R}^d$ and $\mathbf{C} \in \mathbb{R}^{d \times d}$. Here $\mathbb{R}^{d \times d}$ denotes the set of all $d \times d$ symmetric matrices. In this case $m = \frac{1}{2}(d+1)(d+2)$.

Linear Models

Example. Similarly, the family of all quadratic functions can be expressed in vector notation as

$$f(\mathbf{x}; \mathbf{a}, \mathbf{b}, \mathbf{C}) = a + \mathbf{b} \cdot \mathbf{x} + \mathbf{x} \cdot \mathbf{C} \mathbf{x},$$

where $a \in \mathbb{R}$, $\mathbf{b} \in \mathbb{R}^d$ and $\mathbf{C} \in \mathbb{R}^{d \times d}$. Here $\mathbb{R}^{d \times d}$ denotes the set of all $d \times d$ symmetric matrices. In this case $m = \frac{1}{2}(d+1)(d+2)$.

Remark. Dimension m for the family of polynomials in d variables of degree at most ℓ is

$$m = \frac{(d+\ell)!}{d! \ell!} = \frac{(d+1)(d+2) \cdots (d+\ell)}{\ell!}.$$

This grows like d^ℓ as the dimension d grows. This means that these models can become impractical when the dimension d is large. In such cases we can use custom built models rather than general ones.

Model Residuals or Modeling Noise

Recall that given the data $\{(\mathbf{x}_j, y_j)\}_{j=1}^n$ and any model $f(\mathbf{x}; \beta_1, \dots, \beta_m)$, the residual associated with each (\mathbf{x}_j, y_j) is defined by the relation

$$y_j = f(\mathbf{x}_j; \beta_1, \dots, \beta_m) + r_j(\beta_1, \dots, \beta_m).$$

The linear model given by the basis functions $\{f_i(\mathbf{x})\}_{i=1}^m$ is

$$f(\mathbf{x}; \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i f_i(\mathbf{x}),$$

for which the residual $r_j(\beta_1, \dots, \beta_m)$ is given by

$$r_j(\beta_1, \dots, \beta_m) = y_j - \sum_{i=1}^m \beta_i f_i(\mathbf{x}_j).$$

Model Residuals or Modeling Noise

Recall that given the data $\{(\mathbf{x}_j, y_j)\}_{j=1}^n$ and any model $f(\mathbf{x}; \beta_1, \dots, \beta_m)$, the residual associated with each (\mathbf{x}_j, y_j) is defined by the relation

$$y_j = f(\mathbf{x}_j; \beta_1, \dots, \beta_m) + r_j(\beta_1, \dots, \beta_m).$$

The linear model given by the basis functions $\{f_i(\mathbf{x})\}_{i=1}^m$ is

$$f(\mathbf{x}; \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i f_i(\mathbf{x}),$$

for which the residual $r_j(\beta_1, \dots, \beta_m)$ is given by

$$r_j(\beta_1, \dots, \beta_m) = y_j - \sum_{i=1}^m \beta_i f_i(\mathbf{x}_j).$$

The idea is to determine the parameters β_1, \dots, β_m in the statistical model by minimizing the residuals $r_j(\beta_1, \dots, \beta_m)$. In general $m \ll n$ so the residuals may not all vanish.

Linear Models and Residuals: Matrix Notation

This so-called *fitting problem* can be recast in terms of vectors. Define the m -vector β , the n -vectors \mathbf{y} and \mathbf{r} , and the $n \times m$ -matrix \mathbf{F} by

$$\beta = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_m \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} r_1 \\ \vdots \\ r_n \end{pmatrix},$$
$$\mathbf{F} = \begin{pmatrix} f_1(\mathbf{x}_1) & \cdots & f_m(\mathbf{x}_1) \\ \vdots & \vdots & \vdots \\ f_1(\mathbf{x}_n) & \cdots & f_m(\mathbf{x}_n) \end{pmatrix}.$$

Linear Models and Residuals: Matrix Notation

This so-called *fitting problem* can be recast in terms of vectors. Define the m -vector β , the n -vectors \mathbf{y} and \mathbf{r} , and the $n \times m$ -matrix \mathbf{F} by

$$\beta = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_m \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} r_1 \\ \vdots \\ r_n \end{pmatrix},$$

$$\mathbf{F} = \begin{pmatrix} f_1(\mathbf{x}_1) & \cdots & f_m(\mathbf{x}_1) \\ \vdots & \vdots & \vdots \\ f_1(\mathbf{x}_n) & \cdots & f_m(\mathbf{x}_n) \end{pmatrix}.$$

We will assume the matrix \mathbf{F} has rank m . The fitting problem is the problem of finding a value of β that minimizes the "size" of

$$\mathbf{r}(\beta) = \mathbf{y} - \mathbf{F}\beta.$$

But what does "size" mean?

2. Linear Euclidean Least Squares Fitting

A popular notion of the size of a vector is the *Euclidean norm*, which is

$$\|\mathbf{r}(\boldsymbol{\beta})\| = \sqrt{\mathbf{r}(\boldsymbol{\beta})^T \mathbf{r}(\boldsymbol{\beta})} = \sqrt{\sum_{j=1}^n r_j(\beta_1, \dots, \beta_m)^2}.$$

2. Linear Euclidean Least Squares Fitting

A popular notion of the size of a vector is the *Euclidean norm*, which is

$$\|\mathbf{r}(\boldsymbol{\beta})\| = \sqrt{\mathbf{r}(\boldsymbol{\beta})^T \mathbf{r}(\boldsymbol{\beta})} = \sqrt{\sum_{j=1}^n r_j(\beta_1, \dots, \beta_m)^2}.$$

Minimizing $\|\mathbf{r}(\boldsymbol{\beta})\|$ is equivalent to minimizing $\|\mathbf{r}(\boldsymbol{\beta})\|^2$, which is the sum of the “squares” of the residuals.

For linear models $\|\mathbf{r}(\boldsymbol{\beta})\|^2$ is a quadratic function of $\boldsymbol{\beta}$ that is easy to minimize, which is why the method is popular. Specifically, because $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{y} - \mathbf{F}\boldsymbol{\beta}$, we minimize

$$\begin{aligned} q(\boldsymbol{\beta}) &= \frac{1}{2} \|\mathbf{r}(\boldsymbol{\beta})\|^2 = \frac{1}{2} \mathbf{r}(\boldsymbol{\beta})^T \mathbf{r}(\boldsymbol{\beta}) = \frac{1}{2} (\mathbf{y} - \mathbf{F}\boldsymbol{\beta})^T (\mathbf{y} - \mathbf{F}\boldsymbol{\beta}) \\ &= \frac{1}{2} \mathbf{y}^T \mathbf{y} - \boldsymbol{\beta}^T \mathbf{F}^T \mathbf{y} + \frac{1}{2} \boldsymbol{\beta}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\beta}. \end{aligned}$$

2. Linear Euclidean Least Squares Fitting

A popular notion of the size of a vector is the *Euclidean norm*, which is

$$\|\mathbf{r}(\boldsymbol{\beta})\| = \sqrt{\mathbf{r}(\boldsymbol{\beta})^T \mathbf{r}(\boldsymbol{\beta})} = \sqrt{\sum_{j=1}^n r_j(\beta_1, \dots, \beta_m)^2}.$$

Minimizing $\|\mathbf{r}(\boldsymbol{\beta})\|$ is equivalent to minimizing $\|\mathbf{r}(\boldsymbol{\beta})\|^2$, which is the sum of the “squares” of the residuals.

For linear models $\|\mathbf{r}(\boldsymbol{\beta})\|^2$ is a quadratic function of $\boldsymbol{\beta}$ that is easy to minimize, which is why the method is popular. Specifically, because $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{y} - \mathbf{F}\boldsymbol{\beta}$, we minimize

$$\begin{aligned} q(\boldsymbol{\beta}) &= \frac{1}{2} \|\mathbf{r}(\boldsymbol{\beta})\|^2 = \frac{1}{2} \mathbf{r}(\boldsymbol{\beta})^T \mathbf{r}(\boldsymbol{\beta}) = \frac{1}{2} (\mathbf{y} - \mathbf{F}\boldsymbol{\beta})^T (\mathbf{y} - \mathbf{F}\boldsymbol{\beta}) \\ &= \frac{1}{2} \mathbf{y}^T \mathbf{y} - \boldsymbol{\beta}^T \mathbf{F}^T \mathbf{y} + \frac{1}{2} \boldsymbol{\beta}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\beta}. \end{aligned}$$

We will use multivariable calculus to minimize this quadratic function.

The Gradient

Recall that the *gradient* (if it exists) of a real-valued function $q(\beta)$ with respect to the m -vector β is the (unique) m -vector $\partial_{\beta} q(\beta)$ such that

$$\left. \frac{d}{ds} q(\beta + s\gamma) \right|_{s=0} = \gamma^T \partial_{\beta} q(\beta) \quad \text{for every } \gamma \in \mathbb{R}^m.$$

The Gradient

Recall that the *gradient* (if it exists) of a real-valued function $q(\boldsymbol{\beta})$ with respect to the m -vector $\boldsymbol{\beta}$ is the (unique) m -vector $\partial_{\boldsymbol{\beta}} q(\boldsymbol{\beta})$ such that

$$\left. \frac{d}{ds} q(\boldsymbol{\beta} + s\boldsymbol{\gamma}) \right|_{s=0} = \boldsymbol{\gamma}^T \partial_{\boldsymbol{\beta}} q(\boldsymbol{\beta}) \quad \text{for every } \boldsymbol{\gamma} \in \mathbb{R}^m.$$

For the quadratic $q(\boldsymbol{\beta})$ from our least squares problem we have

$$q(\boldsymbol{\beta} + s\boldsymbol{\gamma}) = q(\boldsymbol{\beta}) + s\boldsymbol{\gamma}^T (\mathbf{F}^T \mathbf{F} \boldsymbol{\beta} - \mathbf{F}^T \mathbf{y}) + \frac{1}{2} s^2 \boldsymbol{\gamma}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\gamma}.$$

The Gradient

Recall that the *gradient* (if it exists) of a real-valued function $q(\beta)$ with respect to the m -vector β is the (unique) m -vector $\partial_{\beta} q(\beta)$ such that

$$\left. \frac{d}{ds} q(\beta + s\gamma) \right|_{s=0} = \gamma^T \partial_{\beta} q(\beta) \quad \text{for every } \gamma \in \mathbb{R}^m.$$

For the quadratic $q(\beta)$ from our least squares problem we have

$$q(\beta + s\gamma) = q(\beta) + s\gamma^T (\mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y}) + \frac{1}{2} s^2 \gamma^T \mathbf{F}^T \mathbf{F} \gamma.$$

By differentiating this with respect to s and setting $s = 0$ we obtain

$$\left. \frac{d}{ds} q(\beta + s\gamma) \right|_{s=0} = \gamma^T (\mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y}),$$

from which we read off that the gradient of $q(\beta)$ is

$$\partial_{\beta} q(\beta) = \mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y}.$$

The Hessian

Similarly, the derivative (if it exists) of the vector-valued function $\partial_{\beta} q(\beta)$ with respect to the m -vector β is the $m \times m$ -matrix $\partial_{\beta\beta} q(\beta)$ such that

$$\left. \frac{d}{ds} \partial_{\beta} q(\beta + s\gamma) \right|_{s=0} = \partial_{\beta\beta} q(\beta) \gamma \quad \text{for every } \gamma \in \mathbb{R}^m.$$

The symmetric matrix-valued function $\partial_{\beta\beta} q(\beta)$ is the *Hessian* of $q(\beta)$.

The Hessian

Similarly, the derivative (if it exists) of the vector-valued function $\partial_{\beta} q(\beta)$ with respect to the m -vector β is the $m \times m$ -matrix $\partial_{\beta\beta} q(\beta)$ such that

$$\left. \frac{d}{ds} \partial_{\beta} q(\beta + s\gamma) \right|_{s=0} = \partial_{\beta\beta} q(\beta) \gamma \quad \text{for every } \gamma \in \mathbb{R}^m.$$

The symmetric matrix-valued function $\partial_{\beta\beta} q(\beta)$ is the *Hessian* of $q(\beta)$. For the quadratic $q(\beta)$ from our least squares problem we have

$$\partial_{\beta} q(\beta + s\gamma) = \mathbf{F}^T \mathbf{F}(\beta + s\gamma) - \mathbf{F}^T \mathbf{y} = \partial_{\beta} q(\beta) + s \mathbf{F}^T \mathbf{F} \gamma.$$

The Hessian

Similarly, the derivative (if it exists) of the vector-valued function $\partial_{\beta} q(\beta)$ with respect to the m -vector β is the $m \times m$ -matrix $\partial_{\beta\beta} q(\beta)$ such that

$$\left. \frac{d}{ds} \partial_{\beta} q(\beta + s\gamma) \right|_{s=0} = \partial_{\beta\beta} q(\beta) \gamma \quad \text{for every } \gamma \in \mathbb{R}^m.$$

The symmetric matrix-valued function $\partial_{\beta\beta} q(\beta)$ is the *Hessian* of $q(\beta)$. For the quadratic $q(\beta)$ from our least squares problem we have

$$\partial_{\beta} q(\beta + s\gamma) = \mathbf{F}^T \mathbf{F}(\beta + s\gamma) - \mathbf{F}^T \mathbf{y} = \partial_{\beta} q(\beta) + s \mathbf{F}^T \mathbf{F} \gamma.$$

By differentiating this with respect to s and setting $s = 0$ we obtain

$$\left. \frac{d}{ds} \partial_{\beta} q(\beta + s\gamma) \right|_{s=0} = \left. \frac{d}{ds} (\partial_{\beta} q(\beta) + s \mathbf{F}^T \mathbf{F} \gamma) \right|_{s=0} = \mathbf{F}^T \mathbf{F} \gamma,$$

from which we read off that the Hessian of $q(\beta)$ is

$$\partial_{\beta\beta} q(\beta) = \mathbf{F}^T \mathbf{F}.$$

Convexity and Strict Convexity

We now show that the $m \times m$ -matrix $\mathbf{F}^T \mathbf{F}$ is positive definite. We have

$$\boldsymbol{\gamma}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\gamma} = (\mathbf{F} \boldsymbol{\gamma})^T \mathbf{F} \boldsymbol{\gamma} = \|\mathbf{F} \boldsymbol{\gamma}\|^2 \geq 0 \quad \text{for every } \boldsymbol{\gamma} \in \mathbb{R}^m,$$

whereby $\mathbf{F}^T \mathbf{F}$ is nonnegative definite. It will be positive definite if we can show that

$$\boldsymbol{\gamma}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\gamma} = 0 \quad \implies \quad \boldsymbol{\gamma} = \mathbf{0}.$$

However, because $\boldsymbol{\gamma}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\gamma} = \|\mathbf{F} \boldsymbol{\gamma}\|^2$, it is clear that

$$\boldsymbol{\gamma}^T \mathbf{F}^T \mathbf{F} \boldsymbol{\gamma} = 0 \quad \implies \quad \|\mathbf{F} \boldsymbol{\gamma}\| = 0 \quad \implies \quad \mathbf{F} \boldsymbol{\gamma} = \mathbf{0}.$$

Because \mathbf{F} has rank m , its columns are linearly independent, whereby

$$\mathbf{F} \boldsymbol{\gamma} = \mathbf{0} \quad \implies \quad \boldsymbol{\gamma} = \mathbf{0}.$$

Therefore $\mathbf{F}^T \mathbf{F}$ is positive definite.

The Minimizer

Because $\partial_{\beta\beta} q(\beta) = \mathbf{F}^T \mathbf{F}$ is positive definite, the function $q(\beta)$ is strictly convex, whereby it has a unique global minimizer where its gradient vanishes. Upon setting the gradient of $q(\beta)$ equal to zero, we see that the minimizer satisfies

$$\partial_{\beta} q(\beta) = \mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y} = \mathbf{0} .$$

The Minimizer

Because $\partial_{\beta\beta} q(\beta) = \mathbf{F}^T \mathbf{F}$ is positive definite, the function $q(\beta)$ is strictly convex, whereby it has a unique global minimizer where its gradient vanishes. Upon setting the gradient of $q(\beta)$ equal to zero, we see that the minimizer satisfies

$$\partial_{\beta} q(\beta) = \mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y} = \mathbf{0} .$$

Because the matrix $\mathbf{F}^T \mathbf{F}$ is positive definite, it is invertible. The solution of the above equation is thereby $\beta = \hat{\beta}$ where

$$\hat{\beta} = (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y} .$$

The Minimizer

Because $\partial_{\beta\beta} q(\beta) = \mathbf{F}^T \mathbf{F}$ is positive definite, the function $q(\beta)$ is strictly convex, whereby it has a unique global minimizer where its gradient vanishes. Upon setting the gradient of $q(\beta)$ equal to zero, we see that the minimizer satisfies

$$\partial_{\beta} q(\beta) = \mathbf{F}^T \mathbf{F} \beta - \mathbf{F}^T \mathbf{y} = \mathbf{0}.$$

Because the matrix $\mathbf{F}^T \mathbf{F}$ is positive definite, it is invertible. The solution of the above equation is thereby $\beta = \hat{\beta}$ where

$$\hat{\beta} = (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y}.$$

Remark. In practice you should not compute $(\mathbf{F}^T \mathbf{F})^{-1}$ when $m > 2$. Rather, you should think of the right-hand side above as notation for the solution of the linear algebraic system $\mathbf{F}^T \mathbf{F} \beta = \mathbf{F}^T \mathbf{y}$. All that you need to compute is the solution $\hat{\beta}$ of this system.

The Minimizer

The fact that $\hat{\beta}$ is a global minimizer can be seen from the identity

$$\begin{aligned}q(\beta) &= q(\hat{\beta} + (\beta - \hat{\beta})) \\&= \frac{1}{2} \mathbf{y}^T \mathbf{y} - \frac{1}{2} \hat{\beta}^T \mathbf{F}^T \mathbf{F} \hat{\beta} + \frac{1}{2} (\beta - \hat{\beta})^T \mathbf{F}^T \mathbf{F} (\beta - \hat{\beta}) \\&= q(\hat{\beta}) + \frac{1}{2} (\beta - \hat{\beta})^T \mathbf{F}^T \mathbf{F} (\beta - \hat{\beta}).\end{aligned}$$

Specifically, because $\mathbf{F}^T \mathbf{F}$ is positive definite, we see that:

- $q(\beta) \geq q(\hat{\beta})$ for every $\beta \in \mathbb{R}^m$,
- $q(\beta) = q(\hat{\beta}) \iff \beta = \hat{\beta}$.

Geometric Interpretation: Orthogonal Projections

The least squares fit has a beautiful geometric interpretation with respect to the associated Euclidean inner product

$$(\mathbf{p} \mid \mathbf{q}) = \mathbf{p}^T \mathbf{q}.$$

Define $\hat{\mathbf{r}} = \mathbf{r}(\hat{\boldsymbol{\beta}}) = \mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}$. Observe that

$$\mathbf{y} = \mathbf{F}\hat{\boldsymbol{\beta}} + \hat{\mathbf{r}} = \mathbf{F}(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y} + \hat{\mathbf{r}}.$$

Geometric Interpretation: Orthogonal Projections

The least squares fit has a beautiful geometric interpretation with respect to the associated Euclidean inner product

$$(\mathbf{p} \mid \mathbf{q}) = \mathbf{p}^T \mathbf{q}.$$

Define $\hat{\mathbf{r}} = \mathbf{r}(\hat{\boldsymbol{\beta}}) = \mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}$. Observe that

$$\mathbf{y} = \mathbf{F}\hat{\boldsymbol{\beta}} + \hat{\mathbf{r}} = \mathbf{F}(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y} + \hat{\mathbf{r}}.$$

The matrix $\mathbf{P} = \mathbf{F}(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T$ has the properties

$$\mathbf{P}^2 = \mathbf{P}, \quad \mathbf{P}^T = \mathbf{P}.$$

This means that $\mathbf{P}\mathbf{y}$ is the orthogonal projection of \mathbf{y} onto the subspace of \mathbb{R}^n spanned by the columns of \mathbf{F} .

Geometric Interpretation: Orthogonal Projections

The least squares fit has a beautiful geometric interpretation with respect to the associated Euclidean inner product

$$(\mathbf{p} \mid \mathbf{q}) = \mathbf{p}^T \mathbf{q}.$$

Define $\hat{\mathbf{r}} = \mathbf{r}(\hat{\boldsymbol{\beta}}) = \mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}$. Observe that

$$\mathbf{y} = \mathbf{F}\hat{\boldsymbol{\beta}} + \hat{\mathbf{r}} = \mathbf{F}(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y} + \hat{\mathbf{r}}.$$

The matrix $\mathbf{P} = \mathbf{F}(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T$ has the properties

$$\mathbf{P}^2 = \mathbf{P}, \quad \mathbf{P}^T = \mathbf{P}.$$

This means that $\mathbf{P}\mathbf{y}$ is the orthogonal projection of \mathbf{y} onto the subspace of \mathbb{R}^n spanned by the columns of \mathbf{F} .

Because $\mathbf{F}^T \mathbf{P} = \mathbf{F}^T$ and $\mathbf{y} = \mathbf{P}\mathbf{y} + \hat{\mathbf{r}}$, we see that $\mathbf{F}^T \hat{\mathbf{r}} = 0$. Hence, $\mathbf{y} = \mathbf{F}\hat{\boldsymbol{\beta}} + \hat{\mathbf{r}}$ is the orthogonal decomposition of \mathbf{y} with respect to the subspace spanned by the columns of \mathbf{F} .

A 2-dimensional Example

Example. Least Squares for the affine model $f(t; \alpha, \beta) = \alpha + \beta t$ and data $\{(t_j, y_j)\}_{j=1}^n$. Matrix \mathbf{F} has the form

$$\mathbf{F} = \begin{pmatrix} \mathbf{1} & \mathbf{t} \end{pmatrix}, \quad \text{where } \mathbf{1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}, \quad \mathbf{t} = \begin{pmatrix} t_1 \\ \vdots \\ t_n \end{pmatrix}.$$

Define

$$\bar{t} = \frac{1}{n} \sum_{j=1}^n t_j, \quad \overline{t^2} = \frac{1}{n} \sum_{j=1}^n t_j^2, \quad \sigma_t^2 = \frac{1}{n} \sum_{j=1}^n (t_j - \bar{t})^2,$$

A 2-dimensional Example

Example. Least Squares for the affine model $f(t; \alpha, \beta) = \alpha + \beta t$ and data $\{(t_j, y_j)\}_{j=1}^n$. Matrix \mathbf{F} has the form

$$\mathbf{F} = \begin{pmatrix} \mathbf{1} & \mathbf{t} \end{pmatrix}, \quad \text{where } \mathbf{1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}, \quad \mathbf{t} = \begin{pmatrix} t_1 \\ \vdots \\ t_n \end{pmatrix}.$$

Define

$$\bar{t} = \frac{1}{n} \sum_{j=1}^n t_j, \quad \overline{t^2} = \frac{1}{n} \sum_{j=1}^n t_j^2, \quad \sigma_t^2 = \frac{1}{n} \sum_{j=1}^n (t_j - \bar{t})^2,$$

To obtain:

$$\mathbf{F}^T \mathbf{F} = \begin{pmatrix} \mathbf{1}^T \mathbf{1} & \mathbf{1}^T \mathbf{t} \\ \mathbf{t}^T \mathbf{1} & \mathbf{t}^T \mathbf{t} \end{pmatrix} = n \begin{pmatrix} 1 & \bar{t} \\ \bar{t} & \overline{t^2} \end{pmatrix},$$

$$\det(\mathbf{F}^T \mathbf{F}) = n^2 (\overline{t^2} - \bar{t}^2) = n^2 \sigma_t^2 > 0.$$

Notice that \bar{t} and σ_t^2 are the sample mean and variance of t respectively

The 2-dimensional Example: Explicit Formulas

Then the $\hat{\alpha}$ and $\hat{\beta}$ that give the least squares fit are given by

$$\begin{aligned} \begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix} &= \hat{\boldsymbol{\beta}} = (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{y} = \frac{1}{n} \frac{1}{\sigma_t^2} \begin{pmatrix} \bar{t}^2 & -\bar{t} \\ -\bar{t} & 1 \end{pmatrix} \begin{pmatrix} \mathbf{1}^T \\ \mathbf{t}^T \end{pmatrix} \mathbf{y} \\ &= \frac{1}{\sigma_t^2} \begin{pmatrix} \bar{t}^2 & -\bar{t} \\ -\bar{t} & 1 \end{pmatrix} \begin{pmatrix} \bar{y} \\ \overline{ty} \end{pmatrix} = \frac{1}{\sigma_t^2} \begin{pmatrix} \bar{t}^2 \bar{y} - \bar{t} \overline{ty} \\ \overline{ty} - \bar{t} \bar{y} \end{pmatrix}, \end{aligned}$$

where

$$\bar{y} = \frac{1}{n} \mathbf{1}^T \mathbf{y} = \frac{1}{n} \sum_{j=1}^n y_j, \quad \overline{ty} = \frac{1}{n} \mathbf{t}^T \mathbf{y} = \frac{1}{n} \sum_{j=1}^n y_j t_j.$$

These formulas for $\hat{\alpha}$ and $\hat{\beta}$ can be expressed simply as

$$\hat{\beta} = \frac{\overline{ty} - \bar{y} \bar{t}}{\sigma_t^2}, \quad \hat{\alpha} = \bar{y} - \hat{\beta} \bar{t}.$$

Notice that $\hat{\beta}$ is the ratio of the covariance of y and t to the variance of t .

Least Squares for the General Linear Model

The best fit is thereby

$$\hat{f}(t) = \hat{\alpha} + \hat{\beta}t = \bar{y} + \hat{\beta}(t - \bar{t}) = \bar{y} + \frac{\overline{yt} - \bar{y}\bar{t}}{\sigma_t^2} (t - \bar{t}).$$

Least Squares for the General Linear Model

The best fit is thereby

$$\hat{f}(t) = \hat{\alpha} + \hat{\beta}t = \bar{y} + \hat{\beta}(t - \bar{t}) = \bar{y} + \frac{\overline{yt} - \bar{y}\bar{t}}{\sigma_t^2} (t - \bar{t}).$$

Remark. In the above example we inverted the matrix $\mathbf{F}^T\mathbf{F}$ to obtain $\hat{\beta}$. This was easy because our model had only two parameters in it, so $\mathbf{F}^T\mathbf{F}$ was only 2×2 . The number of parameters m does not have to be too large before this approach becomes slow or unfeasible. However for fairly large m you can obtain $\hat{\beta}$ by using Gaussian elimination or some other direct method to efficiently solve the linear system

$$\mathbf{F}^T\mathbf{F}\beta = \mathbf{F}^T\mathbf{y}.$$

Such methods work because the matrix $\mathbf{F}^T\mathbf{F}$ is positive definite. As we will soon see, this step can be simplified by constructing the basis $\{f_i(t)\}_{i=1}^m$ so that $\mathbf{F}^T\mathbf{F}$ is diagonal.

Further Questions

We have seen how to use least squares to fit linear statistical models with m parameters to data sets containing n pairs when $m \ll n$. Among the questions that arise are the following.

- How does one pick a basis that is well suited to the given data?
- How can one avoid overfitting?
- Do these methods extended to nonlinear statistical models?
- Can one use other notions of smallness of the residual? Maximum Likelihood Estimation.