Lectures on Dynamic Systems and Control

Mohammed Dahleh   Munther A. Dahleh   George Verghese

Department of Electrical Engineering and Computer Science

Massachuasetts Institute of Technology[1]

Chapter 3

Least Squares Solution of $y = \prec A, x \succ$

3.1 Introduction

We turn to a problem that is dual to the overconstrained estimation problems considered so far. Let A denote an array of m vectors, $A = [a_1 | \cdots | a_m]$, where the a_i are vectors from any space on which an inner product is defined. The space is allowed to be infinite dimensional, e.g. the space $\mathcal{L}^2$ of square integrable functions mentioned in Chapter 2 . We are interested in the vector x, *of minimum length*, that satisfy the equation

$$y = \prec A, x \succ \qquad (1a)$$

where we have used the Gram product notation introduced in Chapter 2.

Example 3.1   Let $y[0]$ denote the output at time 0 of a noncausal FIR filter whose input is the sequence $x[k]$, with

$$y[0] = \sum_{i=-N}^{N} h_i x[-i].$$

Describe the set of input values that yield $y[0] = 0$; repeat for $y[0] = 7$. The solution of minimum energy (or RMS value) is the one that minimizes $\sum_{i=-N}^{N} x^2[i]$.

3.2 Constructing all Solutions

When the a_i's are drawn from ordinary (real or complex) Euclidean n-space, with the usual (unweighted) inner product, A is an $n \times m$ matrix of full column rank, and the equation (1a) is simply

$$y = A' x \quad , \qquad (1b)$$

where A' has full row rank. Since the m rows of A' in (1b) are independent, this matrix has m independent columns as well. It follows that the system (1b), which can be read as expressing y in terms of a linear combination of the columns of A' (with weights given by the components of x) has solutions x for any y.

If A' were square and therefore (under our rank assumption) invertible, (1b) would have a unique solution, obtained simply by premultiplying (1b) by the inverse of A'. The closest we come to having an invertible matrix in the non-square case is by invoking the Gram matrix lemma, which tells us that $A'A$ is invertible under our rank assumption. This fact, and inspection of (1b), allow us to explicitly write down one particular solution of (1b), which we denote by $\check{x}$:

$$\check{x} = A\,(A'\,A)^{-1}\,y \qquad (2a)$$

Simple substitution of this expression in (1b) verifies that it is indeed a solution. We shall shortly see that this solution actually has minimum length (norm) among all solutions of (1b).

For the more general equation in (1a), we can establish the existence of a solution by demonstrating that the appropriate generalization of the expression in (2a) does indeed satisfy (1a). For this, pick

$$\check{x} = A\ \prec A, A\succ^{-1}\,y \qquad (2b)$$

It is easy to see that this satisfies (1a), if we use the fact that $\prec A, A\alpha\succ\,=\,\prec A, A\succ \alpha$ for any array α of scalars; in our case α is the $m \times 1$ array $\prec A, A\succ^{-1}\ y$.

Any other x is a solution of (1a) iff it differs from the particular solution above (or any other particular solution) by a solution of the homogeneous equation $\prec A, x\succ\,=0$; the same statement can be made for solutions of (1b). The proof is easy, and presented below for (1b), with x denoting any solution, x_p denoting a particular solution, and x_h denoting a solution of the homogeneous equation:

$$y = A'x_p = A'x \quad \Rightarrow \quad A'\underbrace{(x - x_p)}_{x_h} = 0 \quad \Rightarrow \quad x = x_p + x_h$$

Conversely,

$$y = A'x_p\,, \quad A'x_h = 0 \quad \Rightarrow \quad y = A'\underbrace{(x_p + x_h)}_{x} \quad \Rightarrow \quad x = x_p + x_h.$$

Equations of the form (1a), (1b) commonly arise in situations where x represents a vector of control inputs and y represents a vector of objectives or targets. The problem is then to use some appropriate criterion and/or constraints to narrow down the set of controls.

Example 3.2 Let $m = 1$, so that A' is a single nonzero row, which we shall denote by a'. If $y = 0$, the set of solutions corresponds to vectors x that are orthogonal to the vector a, i.e. to vectors in the orthogonal complement of a, namely in the subspace $\mathcal{R}a^{\perp}(a)$. Use this to costruct all solutions to Example 3.1.

There are several different criteria and constraints that may reasonably be used to select among the different possible solutions. For example, in some problems it may be natural to restrict the components x_i of x to be nonnegative, and to ask for the control that minimizes $\sum s_i x_i$, where s_i represents the cost of control component x_i. This is the prototypical form of what is termed the *linear programming* problem. (You should geometrically characterize the solution to this problem for the case given in the above example.) The general linear programming problem arises in a host of applications.

We shall focus on the problem of determining the solution x of (1a) for which $\|x\|^2 = <x, x>$ is minimized; in the case of (1b), we are looking to minimize $x'x$. For the situation depicted in the above example, the optimum x is immediately seen to be the solution vector that is aligned with a. It can be found by projecting any particular solution of (1b) onto the space spanned by the vector a. (This fact is related to the Cauchy-Schwartz inequality: For x of a specified length, the inner product $<a, x>$ is maximized by aligning x with a, and for specified $<a, x>$ the length of x is minimized by again aligning x with a.) The generalization to $m > 1$ and to the broader setting of (1a) is direct, and is presented next. You should note the similarity to the proof of the orthogonality principle.

3.3   Least Squares Solution

Let x be a particular solution of (1a). Denote by x_A its unique projection onto the range of A (i.e. onto the space spanned by the vectors a_i) and let $x_{A\perp}$ denote the projection onto the space orthogonal to this. Following the same development as in the proof of the orthogonality principle in Lecture 2, we find

$$x_A = A \prec A, A \succ^{-1} \prec A, x \succ \qquad (3a)$$

with $x_{A\perp} = x - x_A$. Now (1a) allows us to make the substitution $y = \prec A, x \succ$ in (3a), so

$$x_A = A \prec A, A \succ^{-1} y \qquad (3b)$$

which is exactly the expression we had for the solution $\check{x}$ that we determined earlier by inspection, see (2b).

Now note from (3b) that x_A is the same for all solutions x, because it is determined entirely by A and y. Hence it is only $x_{A\perp}$ that is varied by varying x. The orthogonality of x_A and $x_{A\perp}$ allows us to write

$$< x, x > = < x_A, x_A > + < x_{A\perp}, x_{A\perp} >$$

so the best we can do as far as minimizing $< x, x >$ is concerned is to make $x_{A\perp} = 0$. In other words, the optimum solution is $x = x_A = \check{x}$.

Example 3.3   For the FIR filter mentioned in Example 3.1, and considering all input sequences $x[k]$ that result in $y[0] = 7$, find the sequence for which $\sum_{i=-N}^{N} x^2[i]$ is minimized. (Work out this example for yourself!)

Example 3.4 Consider a unit mass moving in a straight line under the action of a force $x(t)$, with position at time t given by $p(t)$. Assume $p(0) = 0$, $\dot{p}(0) = 0$, and suppose we wish to have $p(T) = y$ (with no constraint on $\dot{p}(T)$). Then

$$y = p(T) = \int_0^T (T - t)x(t)\, dt = <a(t), x(t)> \tag{4}$$

This is a typical underconstrained problem, with many choices of $x(t)$ for $0 \le t \le T$ that will result in $p(T) = y$. Let us find the solution $x(t)$ for which

$$\int_0^T x^2(t)\, dt = <x(t), x(t)> \tag{5}$$

is minimized. Evaluating the expression in (2a), we find

$$\check{x}(t) = (T - t)y/(T^3/3) \tag{6}$$

How does your solution change if there is the additional constraint that the mass should be brought to rest at time T, so that $\dot{p}(T) = 0$?

We leave you to consider how *weighted* norms can be minimized.

Exercises

Exercise 3.1 Least Square Error Solution We begin with a mini-tutorial on orthogonal and unitary matrices. An *orthogonal matrix* may be defined as a *square* real matrix whose columns are of unit length and mutually orthogonal to each other — i.e., its columns form an *orthonormal* set. It follows quite easily (as you should try and verify for yourself) that:

- the *inverse* of an orthogonal matrix is just its transpose;

- the *rows* of an orthogonal matrix form an orthonormal set as well;

- the usual Euclidean *inner product* of two real vectors v and w, namely the scalar $v'w$, equals the inner product of Uv and Uw, if U is an orthogonal matrix — and therefore the length of v, namely $\sqrt{v'v}$, equals that of Uv.

A *unitary matrix* is similarly defined, except that its entries are allowed to be complex — so its inverse is the *complex conjugate* of its transpose. A fact about orthogonal matrices that turns out to be important in several numerical algorithms is the following: Given a real $m \times n$ matrix A of full column rank, it is possible (in many ways) to find an orthogonal matrix U such that

$$UA = \begin{pmatrix} R \\ 0 \end{pmatrix}$$

where R is a nonsingular, upper-triangular matrix. (If A is *complex*, then we can find a *unitary* matrix U that leads to the same equation.) To see how to compute U in Matlab, read the comments obtained by typing `help qr`; the matrix Q that is referred to in the comments is just U'.

We now turn to the problem of interest. Given a real $m \times n$ matrix A of full column rank, and a real m-vector y, we wish to approximately satisfy the equation $y = Ax$. Specifically, let us choose the vector x to minimize $\|y - Ax\|^2 = (y - Ax)'(y - Ax)$, the squared Euclidean length of the "error" $y - Ax$. By invoking the above results on orthogonal matrices, **show that** (in the notation introduced earlier) the minimizing x is

$$\hat{x} = R^{-1}y_1$$

where y_1 denotes the vector formed from the first n components of Uy. (In practice, we would not bother to find R^{-1} explicitly. Instead, taking advantage of the upper-triangular structure of R, we would solve the system of equations $R\hat{x} = y_1$ by back substitution, starting from the last equation.)

The above way of solving a least-squares problem (proposed by Householder in 1958, but sometimes referred to as Golub's algorithm) is numerically preferable in most cases to solving the "normal equations" in the form $\hat{x} = (A'A)^{-1}A'y$, and is essentially what Matlab does when you write $\hat{x} = A\backslash y$. An (oversimplified!) explanation of the trouble with the normal equation solution is that it implicitly evaluates the product $(R'R)^{-1}R'$, whereas the Householder/Golub method recognizes that this product simply equals R^{-1}, and thereby avoids unnecessary and error prone steps.

Exercise 3.2 Suppose the input sequence $\{u_j\}$ and the output sequence $\{y_j\}$ of a particular system are related by

$$y_k = \sum_{i=1}^{n} h_i u_{k-i}$$

where all quantities are scalar.

(i) Assume we want to have y_n equal to some specified number $\bar{y}$. Determine $u_0, \ldots, u_{n-1}$ so as to achieve this while minimizing $u_0^2 + \ldots + u_{n-1}^2$.

(ii) Suppose now that we are willing to relax our objective of exactly attaining $y_n = \bar{y}$. This leads us to the following modified problem. Determine $u_0, \ldots, u_{n-1}$ so as to minimize

$$r(\bar{y} - y_n)^2 + u_0^2 + \ldots + u_{n-1}^2$$

where r is a positive weighting parameter.

(a) Solve the modified problem.

(b) What do you expect the answer to be in the limiting cases of $r = 0$ and $r = \infty$? Show that your answer in (a) indeed gives you these expected limiting results.

Exercise 3.3 Return to the problem considered in Example 3.4. Suppose that, in addition to requiring $p(T) = y$ for a specified y, we also want $\dot{p}(T) = 0$. In other words, we want to bring the mass *to rest* at the position y at time T. Of all the force functions $x(t)$ that can accomplish this, determine the one that minimizes $< x(t), x(t) >= \int_0^T x^2(t)\, dt$.

Exercise 3.4 (a) Given $y = A'x$, with A' of full row rank, find the solution vector x for which $x'Wx$ is minimum, where $W = L'L$ and L is nonsingular (i.e. where W is Hermitian and positive definite).

(b) A specified current I_0 is to be sent through the fixed voltage source V_0 in the figure. Find what values v_1, v_2, v_3 and v_4 must take so that the total power dissipation in the resistors is minimized.

6.241J / 16.338J Dynamic Systems and Control

Spring 2011