JOURNAL OF COMPUTATIONAL PHYSICS 20, 365–371 (1976)

Notes

Least Square Regression with Multinomials of Restricted Amplitudes*

1. INTRODUCTION

The usual least squares fit to a set of data points yields the minimum standard deviation and the best unbiased estimate of the coefficients of the fit. If one merely wants a compact way of representing a large number of data points, this is a good method to use. However, what is often desired is not the best statistically unbiased fit to the data, but rather the best representation of the function underlying the data; i.e., one wishes to determine whether performance will be improved by a right or left turn of the valve. Thus, the least squares method may give higher order coefficients which are large and of opposing signs, while they should be small and of the same sign. In order to eliminate this effect we will introduce additional information and some bias based on our intuitition or knowledge of the underlying function. Each problem must be treated individually. The appropriate amount of bias to introduce requires careful consideration and is not amenable to a rigorous treatment for a general class of problems. What we present here is the general idea, so that others may adapt it for their own problems. We know of no curve-fitting recipe which can be used blindly.

We will employ two pieces of information which the usual least squares method does not use. First, we include an estimate of the absolute error in each data point, not just the relative error. We use the fact that the function underlying the data misses a typical set of data points by an average of one standard deviation. Our fit must do likewise if it is to have a chance of duplicating the underlying function. To fit the data too closely is just as bad as missing the data by too much. This is justified if we have a large number of points, if the errors have a normal distribution, and if our error estimates are good. Second, we include the knowledge (or perhaps bias) that the underlying function does not have large high-order coefficients. We include these two pieces of information by requiring that the sum of the squares of the coefficients be a minimum, subject to the constraint that the χ^2 be equal to the number of data points.

This method is similar to Ridge Regression [1]. The set of simultaneous equations to be solved for each point along the ridge are the same. However, instead of using

* Work supported in part by a grant from the National Science Foundation.

the characteristics of the function along the ridge to determine the point to be used, we use the simpler criterion that $\chi^2 = N$. This information is usually available for physics-type applications and greatly simplifies the analysis.

2. The Method

Assume that we have N data points, f_k, as a function of the n variables, $z_{1k}, \ldots, z_{nk}$ (where $z_{1k} = 1$, $k = 1, N$). Note that $\{z_{1k}, \ldots, z_{nk}\}$ is, in general, a point in a multi-dimensional space. We wish to fit data with a multinomial

$$F_k = \sum_{j=1}^{n} a_j z_{jk} . \tag{1}$$

In accord with the ideas expressed in the introduction, we minimize the function

$$Q = \chi^2 + C \sum_{j=2}^{n} a_j^2 \tag{2}$$

where

$$\chi^2 = \sum_{k=1}^{N} w_k (F_k - f_k)^2 \tag{3}$$

$$w_k = 1/e_k^2 \tag{4}$$

and e_k is the estimated error in f_k. If $C = 0$, then we have the usual least squares method. The variables should be in standard normal form, but this is generally not appropriate for Physics problems; however, the bias against large high-order coefficients is not lost if the z_j's are simply scaled so that they range from 0 to 1. The $j = 1$ term is excluded from the sum over a_j^2 because a_1 is the constant and we only wish to prevent the other coefficients from becoming large. This form was suggested by that used by Powell for minimizing the discontinuities in the third derivative when curve fitting by splines in one variable [2].

In order to obtain a system of equations for the coefficients, we differentiate Q with respect to a_i and set it equal to zero to obtain

$$\sum_k w_k \left(\sum_j a_j z_{jk} - f_k \right) z_{ik} + C a_i (1 - \delta_{i1}) = 0 \tag{5}$$

or

$$\sum_k \sum_j a_j w_k z_{jk} z_{ik} - \sum_k w_k f_k z_{ik} + C \sum_j a_j \delta'_{ij} = 0 \tag{6}$$

where

$$\delta'_{ij} = \delta_{ij}(1 - \delta_{i1}). \tag{7}$$

We define

$$A_{ij} = \sum_k w_k z_{ik} z_{jk} \tag{8}$$

$$B_i = \sum_k w_k f_k z_{ik} \tag{9}$$

so that

$$\sum_j a_j(A_{ij} + C\delta'_{ij}) - B_i = 0 \qquad i = 1, n. \tag{10}$$

If $C = 0$ this is the usual set of n linear equations for n unknowns, a_j, to be solved simultaneously. We can convert Eq. (10) to this form by defining

$$A'_{ij} = A_{ij} + C\delta'_{ij} \tag{11}$$

That is, simply add C to each of the diagonal elements of the matrix A except the first.

As a bonus this makes the system of equations to be solved less likely to be ill-conditioned, which causes difficulties in obtaining a solution. Now, for any constant C, we can solve the set of equations for the a_i. The method we use is to search for that C which gives $\chi^2 = N$.

If one does not have a sufficient number of variables to achieve the fit to the desired accuracy, e_k, or if the variables have been poorly chosen, it will be impossible to obtain $\chi^2 \leqslant N$. This would generally indicate that either more variables are needed, a more suitable set of variables is needed, or the error term, e_k, has been assumed smaller than the data warrants. It may occur that as C is increased to give $\chi^2 = N$, some of the a_j will go to zero. This seems to be a clear indication that these variables are not needed to achieve a fit to the desired accuracy. Note that the fit with $a_j = 0$ is already a fit without the jth variable unlike the normal polynomial fit which requires another pass when one suspects that one of the variables is unnecessary.

3. Examples of the Method

This method was applied to the problem of fitting a two-dimensional function with random errors. The function fitted was

$$f_k = g(x_k, y_k) + \Delta_k \tag{12}$$

where (x_k , y_k) is a compact notation for the vector

$$(1, x_k , y_k , x_k^2, x_k y_k , y_k^2, x_k^2 y_k , x_k y_k^2, x_k^2 y_k^2) = (z_{1k} ,..., z_{9k})$$

and where

$$g(x, y) = \sum_{i=1}^{9} z_i = 1 + x + y + x^2 + xy + y^2 + x^2 y + xy^2 + x^2 y^2, \quad (13)$$

and Δ_k was a random variable chosen from a set of random numbers generated with a gaussian distribution with mean 0 and a standard deviation of $\sigma = 0.1$. We used an evenly spaced rectangular mesh of 16 points with a range from 0 to 1 in both x and y. While one of the major problems of curve fitting is to choose the proper functional form to fit the data, here we assumed the correct functional form; i.e., we fit f_k with

$$F(x, y) = a_1 + a_2 x + a_3 y + a_4 x^2 + a_5 xy + a_6 y^2 + a_7 x^2 y + a_8 xy^2 + a_9 x^2 y^2. \quad (14)$$

We generated 4 different sets of random fluctuations, Δ_k, and fit the function in Eq. (12) for each data set. The resulting coefficients for all 4 data sets are given in Table I for both the biased and unbiased fits. For the biased fits we assumed $e_k = \sigma = 0.1$. Table I also lists the constant, C, χ_0^2, a measure of the deviation from the data points,

$$\chi_0^2 = \sum_{k=1}^{16} (f_k - F(x_k , y_k))^2 / e_k^2, \quad (15)$$

and χ_1, the deviation from the underlying function,

$$\chi_1^2 = \sum_{k=1}^{16} (g(x_k , y_k) - F(x_k , y_k))^2 / e_k^2. \quad (16)$$

We were unable to obtain a biased fit for the fourth data set because the χ_0^2 for the unbiased fit was already greater than 16. For the first three data sets, however, it is seen that the coefficients from the biased fits are dramatically superior to the unbiased coefficients (a perfect fit would have given $a_i = 1$, $i = 1, 9$.). Also note that the coefficients are reproducible from one data set to another when biasing is used, but not the unbiased fits. The χ_1^2's for the biased fits are smaller than those for the unbiased fits, again indicating better agreement with the underlying function. The most dramatic improvement is in the sum of squares of the

TABLE I

Comparison of Coefficients in Biased and Unbiased ($C = 0$) Fits for 4 Equivalent Sets of Data Points[a]

Data set	1	2	3	4	1	2	3
C	0	0	0	0	31.2	21.1	25.1
χ_0^2	2.1	8.5	8.0	17.9	16.0	16.0	16.0
χ_1^2	8.2	12.1	5.0	10.9	2.9	10.8	2.9
a_1	0.86	1.22	.84	1.01	1.03	1.17	1.02
a_2	1.70	0.19	1.56	1.41	0.97	0.78	0.99
a_3	1.46	0.43	1.60	1.00	0.91	0.88	1.03
a_4	0.42	1.67	0.60	0.62	1.03	1.13	1.04
a_5	−3.80	3.36	−1.92	−2.22	1.01	1.03	1.02
a_6	0.64	1.41	0.60	1.08	1.04	0.93	1.00
a_7	5.60	−0.64	3.71	4.55	0.98	1.11	.99
a_8	5.42	−1.34	3.35	4.30	0.99	0.91	0.92
a_9	−3.24	2.79	−1.33	−2.91	0.91	0.97	0.87
E	82.9	18.6	27.8	49.5	0.02	0.14	0.03

[a] χ_0^2 is the deviation from the data. χ_1^2 is the deviation from the underlying function. E is the sum of the squares of the errors of the a_i.

TABLE II

Comparison of Coefficients in Unbiased ($C = 0$) and Biased Fits to $g(x, y) = x + y$ [a]

C	0	31.6
χ_0^2	8.0	16.0
χ_1^2	5.0	4.9
a_1	−0.03	0.01
a_2	1.11	0.90
a_3	1.12	0.91
a_4	−0.08	0.10
a_5	−0.58	0.14
a_6	−0.08	0.09
a_7	0.54	−0.02
a_8	0.47	−0.05
a_9	−0.47	−0.09
E	1.11	0.07

[a] χ_0^2 is the deviation from the data. χ_1^2 is the deviation from the underlying function, g. E is the sum of the squares of the errors of the a_i.

errors of the coefficients. Figure 1 illustrates the relationships between the underlying function, the data points used for one of the data sets, the unbiased fit and the biased fit for one of the data sets. Even with the deviations from the underlying function magnified by a factor of 10, the biased function maintains the same curvature in each of the four cross-sections; the unbiased fit, on the other hand, changes curvature twice in going from $x = 0$ to $x = 1$, under this magnification.

We also generated one data set from an underlying function

$$g(x, y) = x + y$$

with $\sigma = 0.02$. This function is presumably not as well suited for the biased method since we used the same 9-term $F(x, y)$ given by Eq. (14) to fit the data and all the coefficients should not be the same size. The coefficients are given in Table II. Here the superiority of the biased method is not quite so dramatic, but it definitely improves the fit.

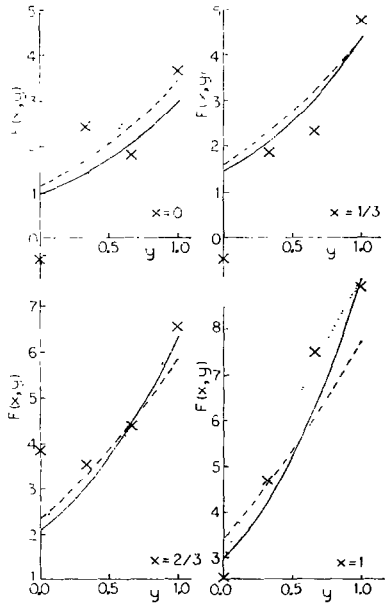

FIG. 1. Relationships between the underlying function, the data points used, the unbiased fit, and the biased fit for one of the data sets generated in four cross-sections of the two-dimensional space. The deviations of the last three functions from the underlying function have been magnified by a factor of 10 for illustrative purposes. Solid curve, underlying function, $g(x, y)$ in Eq. (13); x's, $g(x, y) + 10$[data points-$g(x, y)$]; dotted line, $g(x, y) + 10$ [unbiased fit-$g(x, y)$]; dashed line, $g(x, y) + 10$[biased fit-$g(x, y)$].

4. Conclusions

The method for smoothing data which has been presented here is one way of putting into the least squares fitting routine something which we with our prejudices "know" in addition to the actual data; i.e., that the coefficients of the fit which we want are not large and of opposing signs, but are rather small and to an extent have the same signs. There are no doubt many variations of biasing in which more or less of the above prejudice is used. For example, one might include only coefficients above a sertain order in the sum in Eq. (2) or weight the coefficients above a certain order in the sum in Eq. (2) or weight the coefficients with j or j^2 in the sum so as to further discriminate against high powers.

The biased fit will, of course, not be as close to the data points as the unbiased fit. It does not like sharp peaks and sudden changes (which might just be random errors in the data). Only if the data demands it will the fit reluctantly indicate these. This is useful when one wishes to avoid the possibility of spurious peaks in the fit, but one should be cautious in investigating peaks with biasing.

It is impossible with a simple example to convey the frustration involved in fitting real data. Low-order polynomials miss the data giving too large a χ^2, while high-order polynmials insist on looping through the data and have huge coefficients of opposing signs forcing one to use double and then quadruple precision arithmetic to evaluate a polynomial fit to data with only 3 significant figures. Curve fitting is still partially an art. There is no substitute for choosing the correct functional form. However, by sacrificing a small amount of closeness of fit to the data, the method presented here eliminates the large coefficients. As a bonus, it makes the set of equations involved less ill-conditioned. Also, the biased fit may be a better representation of the function underlying the data. In particular, it may give a much better representation of the derivative of the underlying function, as shown in Fig. 1.

References

1. A. E. Hoerl, *Chem. Eng. Prog.* **58**, 54 (March 1962); D. W. Marquardt and R. D. Snee, "Ridge Regression in Practice," American Statistician **29** (1975), 3.
2. M. J. D. Powell, Curve Fitting by Splines in One Variable, *in* "Numerical Approximation to Functions and Data" (J. G. Hayes, Ed.), p. 65, Athlone Press, London (1970).

Received: June 9, 1975.

R. L. Coldwell
M. A. Pokrant

Department of Physics and Astronomy
University of Florida
Gainsville, Florida 32611