

A Quasi-Newton Procedure for Identifying Pilot-Related Parameters
of the Optimal Control Model

by

William H. Levison
Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, MA 02238

ABSTRACT

Progress is reported in the development and application of a quasi-Newton gradient search procedure for identifying independent pilot-related parameters of the optimal control model for pilot/vehicle systems. The computational efficiency of the scheme originally implemented by Lancraft and Kleinman has been improved. A sensitivity analysis procedure is described that allows one to determine (1) whether or not a given model parameter is required to match a specific experimental result, and (2) which experimentally-induced parameter changes are "significant"; i.e., required to account for behavioral and performance differences. Application of the identification scheme to training effects in a manual control task is described.

INTRODUCTION

Considerable effort has been expended over nearly four decades to develop mathematical models for predicting and diagnosing human operator response behavior in closed-loop control tasks. Predictive models are desired for the general purpose of extrapolating knowledge, gained from man-in-the-loop studies, to tasks in which experimental data have not yet been obtained. Diagnostic models, on the other hand, are intended to help quantify and interpret the effects of stress, and other aspects of the task environment, on operator response capabilities.

Although the same model form may be used for both prediction and diagnosis, the treatment of independent model parameters is different. For predictive applications one desires a set of independent model parameters that are either constant or selectable on the basis of well-defined adjustment rules. For diagnostic applications -- particularly when the influence of environmental or task parameters is not well understood -- independent model parameters will often be adjusted to provide the best match to the experimental data, with little constraint on parameter values. For this type of model usage, a well-defined procedure is required to uniquely identify (i.e., quantify) independent model parameters from the experimental data, and to indicate the reliability of the identified parameter values.

Particular impetus for developing reliability metrics for identified parameters has been provided by a study of the effects of learning on the human controller's response strategy [1]. As shown later in this paper, parameter values undergo rather large changes during training, and a determination as to which of these changes are statistically meaningful can provide guidance for the development of models for learning behavior.

This paper describes a procedure for identifying and testing parameters of the "optimal control model" (OCM) for the human operator in steady-state control tasks. Typical independent -- or "pilot-related" -- parameters to be identified from laboratory tracking data are time delay, motor time constant, (equivalently, a "cost" penalty on rate-of-change of control), motor noise covariance, and an observation noise covariance for each perceptual input variable used by the operator. Readers unfamiliar with this model are directed to the recent review article by Baron and Levison [2] and to the references cited therein.

REVIEW OF THE QUASI-NEWTON IDENTIFICATION PROCEDURE

In this section we first review the process of adjusting independent model parameters to provide a best match to the experimental data. We then offer a technique for deriving an analytic approximation to the sensitivity of the matching error to perturbations in these parameters. Finally, certain implementation details are discussed. In the interest of conciseness, only major results are presented here. Derivations have been reported in more detail in [1].

The Basic Minimization Procedure

Consider the task of adjusting model parameters to minimize a scalar matching error $J = \underline{e}' \underline{W} \underline{e}$, where each element e_i of the column vector \underline{e} is the difference between the i th measured data point and the corresponding model prediction, and each element w_i of the diagonal matrix \underline{W} is a weighting coefficient. In a particular application, the matching error J will correspond to a particular choice of parameter values \underline{p} . The objective of the search procedure is to find a new parameter set $\underline{p} + \Delta \underline{p}$ such that J is minimized.

To implement the search scheme, we initially assume that model predictions (and, therefore, prediction errors) vary linearly with model parameters. Thus, $\Delta \underline{e} = \underline{Q}' \Delta \underline{p}$, where

$$q(i,j) = \partial e_i / \partial p_j \quad (1)$$

Solving for minimum J as a function of \underline{p} , we obtain

$$\Delta \underline{p} = -[\underline{Q} \underline{W} \underline{Q}']^{-1} \underline{Q} \underline{W} \underline{e} \quad (2)$$

Now, since model input/output relationships are seldom totally linear, two or more iterations of the procedure are required until some convergence criteria are satisfied. In some cases, the parameter change computed as shown in Eq(2) will yield a scalar matching error greater than the starting value. Therefore, it is often useful to augment the minimization procedure described above with a line-search scheme to optimize the magnitude of p .

Sensitivity Analysis

An indication of parameter estimation reliability can often be obtained through sensitivity analysis relating changes in the scalar matching error to perturbations in model parameters. In general, estimates of parameters that have a high impact on matching error can be considered more reliable than estimates of parameters having a smaller impact.

If model predictions are linear in the parameters, we may analytically derive the sensitivity of the scalar modeling error to perturbations in model parameters about the optimal (best-matching) set. One may compute the sensitivity to a given parameter with the remaining model parameters held fixed, or with remaining parameters reoptimized. We shall compute sensitivity according to the latter definition because, by allowing tradeoffs among parameters in terms of minimizing matching error, it provides a more stringent reliability measure.

The sensitivity of the matching error J to a change in a single parameter p_i is

$$J = \underline{v}' \underline{Q} \underline{W} \underline{Q}' \underline{v} (\Delta p_i)^2 \quad (3)$$

where J is the increment in J about its minimum value, p_i the change in parameter p_i about its optimum, and \underline{v} is a column vector that has unity value for the i th element and values for remaining elements given by the following expression:

$$\underline{v}_r = -[\underline{Q}_r \underline{W} \underline{Q}'_r]^{-1} \underline{Q}_r \underline{W} \underline{q}_i \quad (4)$$

where $\underline{q}_i = \text{col}(q_{i,1}, q_{i,2}, \dots)$, and the subscript "r" indicates vectors and matrices with omission of rows and columns corresponding to the i th model parameter. (See [1] for a derivation of this result.) The change in matching error, therefore, varies as the square of the change in parameter value, given the underlying assumption of linearity between model parameters and model predictions.

Implementation of Manual Control Studies

Application of the QN method for analysis of human operator performance in continuous control tasks has been reported by Lancraft and Kleinman [3]. Described below is a revised implementation that was used to perform the model analysis described later in this paper.

Two criteria must be defined in order to apply the identification procedure: (1) a definition of a scalar matching error to be minimized by the QN scheme, and (2) convergence criteria to determine when the minimum modeling error has been approached sufficiently closely to justify termination of the minimization procedure.

Matching error is similar to that used by Lancraft and Kleinman:

$$J = \frac{1}{N_1} \sum_{i=1}^{N_1} \left(\frac{G_i - \hat{G}_i}{\sigma_{G_i}} \right)^2 + \frac{1}{N_2} \sum_{i=1}^{N_2} \left(\frac{P_i - \hat{P}_i}{\sigma_{P_i}} \right)^2 \quad (5)$$
$$+ \frac{1}{N_3} \sum_{i=1}^{N_3} \left(\frac{R_i - \hat{R}_i}{\sigma_{R_i}} \right)^2 + \frac{1}{N_4} \sum_{i=1}^{N_4} \left(\frac{S_i - \hat{S}_i}{\sigma_{S_i}} \right)^2$$

where N is the number of valid measurements in the j th measurement group; $G, P,$ are the gain (dB) and phase shift (degrees) of the i th describing function point to be matched; R is the corresponding control-stick "remnant" measurement (dB); and S is the i th variance score to be matched (units different for different tracking variables). indicates standard deviation of an experimental data point, and the symbol " $\hat{}$ " ("hat") indicates a model prediction.

Inclusion of the experimental deviations in the scalar modeling error allows each error component to be weighted inversely by the reliability of the data. To prevent the matching criterion from giving excessive weights to variables that have very low experimental variability, the following minimum standard deviations are imposed: 0.5 dB for magnitude and remnant, 3 degrees for phase, and 5% for the ensemble mean for variance scores.

Weighting inversely by standard deviation also converts each error term into a dimensionless number, thereby allowing accumulation of matching errors into a single metric. Thus, the matching error defined in Eq(5) approximates the average number of standard deviations of mismatch. A numerical score of $J=4$ reflects an average modeling error of 1 standard deviation (i.e., an average error of unity per measurement group).

The minimization procedure is terminated when the following conditions jointly obtain for two successive iterations: (1) reduction of the matching error by less than 0.5%, and (2) changes in all identified parameters by less than 2%. The first criterion is based on the fact that the sensitivity of matching error to small perturbations of model

parameters is relatively low in the vicinity of the minimum (a consequence of the quadratic matching error). The second criterion prevents termination resulting from a compensating "overshoot"; i.e., a situation in which successive estimates of one or more parameters bound the optimal values in such a way as to yield essentially the same modeling error.

A number of modifications have been made to the original implementation in order to improve computational efficiency. First, the search is performed on the logarithms of the parameters. This transformation modestly increases the degree of linearity between model parameters and model outputs, and it prevents the assignment of out-of-bounds (i.e., negative) values to parameters during the course of the search. Second, in order to minimize numerical difficulties with inversion of the expression QWQ' , we omit from the search procedure (i.e., keep fixed), at a given iteration, any parameter having a negligible influence on the matching error. In addition, to reduce the chance of convergence to a local minimum appreciably removed from the global minimum, an individual parameter is allowed to undergo no more than a ten-fold increase or decrease from one iteration to the next.

Finally, a binary section scheme is employed to prevent divergence of the QN scheme due to nonlinear relationships between model inputs and outputs. If necessary, binary section is repeated until (1) matching error is reduced from one iteration to the next, or (2) until four attempts fail to reduce matching error, at which point the minimization scheme is terminated. Further details regarding implementation are documented by Levison [1].

As is true with any numerical search procedure, the probability of convergence to a global minimum is enhanced by the selection of an initial set of model parameters that are close to the optimal set. The following rules for initializing model parameters appears to yield good results with the QN procedure: (1) cost of control rate such that motor time constant = 0.1 seconds; (2) time delay = 0.2 seconds; (3) observation noise covariance to achieve a noise/signal ratio of -20 dB for each perceptual variable assumed to be utilized by the operator; and (4) motor noise covariance to achieve a noise signal ratio of -50 dB, normalized with respect to control-rate variance.

SIGNIFICANCE TESTING

In the following discussion we assume that the data base being subjected to model analysis reflects a significant difference in human operator response behavior, as determined by some standard quantitative test for significance. We then wish to test the hypothesis that the various data sets can be modeled by the same set of model parameters. Failure to support this null hypothesis indicates that parameter differences are also significant.

A cross-comparison scheme was developed and tested against data obtained in manual control studies. In general, this method may be employed to provide a qualitative significance test on parameter differences obtained from modeling the results of two experimental conditions. This method employs an empirical sensitivity test as described below.

Assume that model parameters have been identified from two data sets corresponding to, say, the "baseline" and "test" experimental conditions; our task is now to test the null hypothesis that a single set of model parameters provides a near-optimal match to the baseline and test data. To perform this test, we first identify the following three sets of pilot parameters: (1) the set that best matches the baseline data, (2) the set that best matches the test data, and (3) the set that provides the best joint match to the baseline and test data. For convenience, we shall refer to the parameters identified in step 3 as the "average parameter set".

We next compute the following four matching errors:

$J(B,B)$ = matching error obtained from baseline data, using parameters identified from baseline data (i.e., best match to baseline data).

$J(B,A)$ = matching error obtained from baseline data, using average parameter set.

$J(T,T)$ = best match to test data.

$J(T,A)$ = matching error obtained from test data, using average parameter set.

Finally, we compute the following "matching error ratios": $MER(B) = J(B,A)/J(B,B)$, $MER(T) = J(T,A)/J(T,T)$, and, if we wish to reduce the results to a single number, the average of these two error ratios. In a qualitative sense, the greater the matching error ratios, the more significant are the differences between the parameters identified for the baseline and test conditions.

As shown by Levison [1], a good approximation to the joint match to multiple data sets can be obtained by simply matching the average data. Thus, to obtain the "average parameter set", one would first obtain a point-by-point ensemble average of the (reduced) baseline and test data, and then identify parameters to match the average data set. This procedure is valid if the same task description applies to the two experimental conditions; i.e., if both tasks can be modeled identically except for quantitative differences in pilot-related parameters. Experiments designed to explore training effects, environmental stress, or interference from other concurrent tasks often meet this restriction.

In addition to providing a collective test of the entire parameter set, this scheme may also be used to test a single parameter or a subset of parameters. Suppose, for example, one wishes to test apparent differences

in the time delay parameter. The matching errors $J(B,B)$ and $J(T,T)$ would be computed as described above. The errors $J(B,A)$ and $J(T,A)$, however, would be computed with only the time delay parameter fixed at its "average" value; remaining parameters would be re-optimized.

APPLICATION TO STUDIES OF HUMAN OPERATOR PERFORMANCE

Two applications of the cross-comparison scheme for significance testing are illustrated below. First, data from rate- and acceleration-control systems are analyzed to determine the degree of parameterization required in each case. Second, we analyze the effects of training on operator response behavior.

Parameterization Requirements

The QN analysis methodology described above was applied to data obtained from two manual control studies: one utilizing a rate-control system [4], and one using approximate acceleration-control dynamics [5]. In both studies, a pseudo-random forcing function was applied in parallel with the operator's control input, and subjects were trained to near asymptotic levels of performance. The data bases subjected to model analysis were obtained by averaging performance measures from three subjects for the first study, and from eight subjects for the second.

The following five independent model parameters were identified in each case: (1) observation noise on error, (2) observation noise on error rate, (3) pseudo motor noise, (4) time delay, and (5) relative cost of control rate (equivalently, motor time constant). Identification was repeated for each data base with time delay, pseudo motor noise, and rate observation noise omitted individually from the analysis.* When any one parameter was omitted, remaining parameters were re-optimized to yield minimum modeling error.

Matching error ratios were computed by normalizing the scalar modeling error obtained with a parameter omitted, to the modeling error obtained with all five parameters identified. The matching error ratios presented in Table 1a indicate that all three parameters tested were required to parameterize the data obtained from the rate-control system. That is, with any single parameter omitted, the matching error increased by a factor of three or more. Time delay and rate observation noise were also required to match the acceleration-control data, but pseudo motor noise proved to be an extraneous parameter (matching error ratio of 1.02) for this data set.

* The mathematical structure of the model requires finite, non-zero values for cost of control rate and (for these tasks) for observation noise on error. Therefore, model analysis was not performed with these parameters omitted.

Table 1. Model Parameterization Requirements

Model Parameter	Rate Control	Acceleration Control
-----------------	--------------	----------------------

a) Effect of Omitting Parameter on Matching Error Ratio

Rate Observation Noise	8.6	8.3
Motor Noise	3.6	1.0
Time Delay	65.3	10.7

b) Inverse Sensitivity

Displacement Obs. Noise	1.8	33.3
Rate Observation Noise	2.1	5.9
Motor Noise	4.7	62.2
Time Delay	0.4	2.3
Cost of Control Rate	2.8	5.3

Table 1b contains the analytic inverse sensitivity computations for each of the five parameters identified in the initial analysis for each data base. These measures, which indicate the decilog change required to increase matching error by 4 units, were computed analytically during the QN search as part of the parameter reduction procedure described earlier.

The analytic sensitivity predictions correlate well with the empirical matching error ratios shown in Table 1a. For a given control system, matching error ratio (a direct measure of sensitivity) varies inversely with predicted inverse sensitivity. In particular, especially large inverse sensitivity is shown for the one parameter (motor noise, acceleration control task) that is considered extraneous. Therefore, the analytic sensitivity prediction provides guidance to required model parameterization.

Three conclusions can be drawn from this illustration. First, the gradient search technique in general, and the QN identification scheme in particular, has the intrinsic capability to identify time delay, motor noise, and rate observation noise -- a capability that has not been demonstrated by maximum likelihood schemes [6]. Second, analytic sensitivity computations performed as part of the QN search procedure provide an indication of the required parameterization. Finally, the ability to identify a particular model parameter will, in general, depend on the specifics of the experimental data base.

A Study of Training-Related Performance Differences

The model analysis scheme described in this paper was used in a recent study to quantify and interpret the effects of training on human operator performance [1]. The experimental data base was obtained from an earlier study which explored the effects of delayed motion cuing on roll-axis tracking performance. Of interest here are pre-transition performance measures obtained from subjects initially trained fixed-base.

Subjects were required to maintain simulated wings-level attitude in a single-axis laboratory tracking task. Vehicle dynamics were representative of a high-performance fighter aircraft in the roll axis, and a zero-mean gust environment was simulated. Except for a brief familiarization period, all training and data trials were conducted with the external forcing function and were digitally recorded for subsequent analysis and modeling. Details of the experiment have been reported by Levison, Lancraft, and Junker, [7].

Frequency response measures are shown in Figure 1 for a single subject very early in training ("Early Training") and for the final pre-transition training session ("Late Training"). This training interval represented about 70 experimental trials. Training induced the following changes in response behavior: (1) an increase in amplitude ratio ("gain") at all frequencies, (2) a decrease in high-frequency phase lag, and (3) a reshaping of the control-stick remnant spectrum to yield decreased remnant power at low frequencies and increased remnant at high frequencies. RMS tracking error (not shown) decreased by almost a factor of two over the course of this training interval.

These gain and phase-shift changes are consistent with improved tracking efficiency. While not obvious, training-related changes in remnant are also indicative of improved tracking efficiency and are consistent with the hypothesis (borne out by model analysis) that training leads to decreased response variability and increased man/machine response bandwidth.

Table 2 shows pilot-related model parameters for two test subjects. Parameters are shown for an average of 2-4 trials very early in training and for the average of the final four trials. (The smooth curves shown in

Table 2. Effects of Training on Pilot-Related Model Parameters

State of Training	Subject	Pilot Parameter				
		P_{ye}	$P_{y\dot{e}}$	P_u	τ	T_n
Early	CP	-5.3	-18.6	-28.2	.230	.343
Late		-21.6	-16.4	-29.3	.162	.169
Early	TB	-11.0	-15.9	-70.1	.198	.162
Late		-21.2	-17.4	-56.9	.219	.121

P_{ye} = error observation noise/signal ratio, dB

$P_{y\dot{e}}$ = error rate observation noise/signal ratio, dB

P_u = motor noise/signal ratio, dB

T_n = motor time constant, seconds

τ = time delay, seconds

Figure 1 are model predictions obtained with the parameter sets shown for Subject CP.) To be consistent with previous publications, the relative weighting coefficient for control rate is shown as an equivalent motor time constant [1], and noise covariances are presented as noise/signal ratios.

The following effects of training are noted: (a) a substantial reduction in the observation noise associated with perception of tracking error, (b) a substantial reduction in the motor time constant, (c) a sizeable decrease in time delay for one subject, and (d) an apparently large increase in motor noise for the other subject. Surprisingly, training had small and inconsistent effects on observation noise related to utilization of error rate information.

The cross-comparison significance test was applied to determine which of the identified parameter changes reflected real differences in operator behavior. Tests were performed for the following sets of parameters: (a)




Figure 1. Effects of Training on Frequency Response
Subject CP, average of 4 trials.

the entire set, (b) observation noise parameters as a group, (c) motor time constant, and (d) time delay. Modeling error ratios were computed separately for subjects CP and TB.

Table 3 shows that, taken as a whole, changes in pilot-related model parameters were highly significant. Average model parameters yielded matching errors that were from about 8 to 20 times as great as those obtained with the optimal parameter sets. This result is not surprising, given the substantial training-related changes in operator response behavior shown in Figure 1.

Training-related differences in both the motor time constant parameter and the noise parameter group were important. Differences associated with motor time constant were more significant in the sense that error ratios for this grouping were about 50% higher than ratios associated with the noise parameters.

Fixing the time delay by itself yielded error ratios only slightly greater than unity. A test performed on the large training-related change

Table 3. Test of Model Parameter Differences Due to Training

Parameter Set Tested	Modeling Error Ratio	
	Subject CP	Subject TB
All Parameters	18.5	7.8
All Noises	3.2	2.0
Motor Time Constant	5.0	3.6
Time Delay	1.1	1.0

in motor noise found for subject TB also yielded negligible change in modeling error ratio. Thus, training-related effects on identified changes in motor time constant and observation noise appear to reflect true changes in operator response capabilities, whereas changes in time delay and motor noise are more likely to reflect (for these specific data sets) problems in parameter identification.

DISCUSSION

The requirement for a given parameter to be included in the identified set, and the ease and precision with which the parameter can be identified, are not intrinsic properties of the model parameter in question. Rather, these factors depend partly on the details of the task structure and of the analysis procedures. For example, we showed above that motor noise was required to obtain minimum matching error for one task but not for another.

Parameterization and identifiability will also depend strongly on the experimental measurement set used to define the matching error, and on the set of model parameters being identified. For example, sensitivity analysis performed in other studies [8] suggests that omission of the remnant spectrum from the measurement set would lead to considerable difficulty in distinguishing among the various observation noise sources (and possibly in distinguishing observation noise from time delay). Similarly, if one were to attempt to identify cost weightings for all state variables, along with the independent pilot parameters considered in this paper, overparameterization might well impede identification of one or more parameters.

Caution should be exercised when interpreting the training-related changes in model parameters reported above. As in previous studies, model analysis was based, in part, on the assumption that the subject has a near-perfect internal representation of the task environment (plant dynamics, input spectrum, etc.). While this assumption is appropriate for well-trained subjects tracking with relatively low-order plants, it is less likely to apply to subjects early in training.

More comprehensive analysis of the data base suggests that training-related changes in motor time constant do not reflect differences in motor response capabilities, but other kinds of response limitations not adequately reflected by the model as applied to this study (Levison, 1981). For example, the large motor time constant found early in training may reflect a cautious control strategy (i.e., low pilot gain) arising from the subject's uncertainty with regard to the dynamical response characteristics of the controlled element. Further research is contemplated in this area.

ACKNOWLEDGEMENT

This work was supported by the Air Force Office of Scientific Research under Contract No. F49620-8-C-0073.

REFERENCES

1. Levison, W.H., "Effects of Whole-Body Motion Simulation on Flight Skill Development", Report No. 4645, Bolt Beranek and Newman Inc., Cambridge, MA, 1981 (in preparation).
2. Baron, and W.H. Levison, "The Optimal Control Model: Status and Future Directions", Proceedings of the 1980 International Conference on Cybernetics and Society, Cambridge, MA, Oct. 8-10, 1980.
3. Lancraft, R.E. and D.L. Kleinman, "On the Identification of Parameters in the Optimal Control Model", Proceedings of the Fifteenth Annual Conference on Manual Control, Wright State University, Dayton, OH, March 1979.
4. Levison, W.H., "The Effects of Display Gain and Signal Bandwidth on Human Controller Remnant," AMRL-TR-70-93, Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, OH, March 1971.
5. Levison, W.H., "Modeling the Pilot's Use of Motion Cues During Transient Aircraft Maneuvers", Report No. 4312, Bolt Beranek and Newman Inc., Cambridge, MA, March 1980.
6. Levison, W.H., R.E. Lancraft, and A.M. Junker, "The Effects of Simulator Delays on Performance and Learning in a Roll-Axis Tracking

Task", paper presented at the Fifteenth Annual Conference on Manual Control, Wright State University, Dayton, OH, March 1979.

7. Phatak, A., H. Weinert, I. Segall, and C. Day, "Identification of a Modified Optimal Control Model for the Human Operator", Automatica, Vol. 12, pp. 31-41, 1976.
8. Zacharias, G.L. and W.H. Levison, "A Performance Analyzer for Identifying Changes in Human Operator Tracking Strategies", ARML-TR-79-17, Aerospace Medical Research Laboratory, Wright-Patterson Air Force Base, March 1979.