Pre-Estimation for Better Initial Guesses

Steven R. Shaw, *Member, IEEE*, Marc Keppler, and Steven B. Leeb, *Senior Member, IEEE*

***Abstract*—Fast, assured convergence of iterative identification methods is important in many on-line applications, including systems that exploit indirect measurements, diagnostics, and model-based control. The transient response of a system can often be mapped directly to an estimate of the system model parameters. This initial estimate, or pre-estimate, can be evaluated in fixed time and may greatly accelerate subsequent iterative system identification procedures. A pre-estimation structure is developed and demonstrated with examples including single-tone frequency estimation and induction motor parameter identification.**

***Index Terms*—Estimation, indirect measurement, modeling, system identification.**

I. Introduction

NEW MEASUREMENT and diagnostic techniques depend increasingly on fast, reliable parameter identification routines. Unfortunately, without an initial guess, most iterative identification routines converge slowly, converge to local minima, or fail. Convergence problems are tolerable when a skilled operator can supply preliminary estimates for the initial guess and restart a failed identification procedure. In contrast, identification methods embedded in an instrument, "soft-sensor" [1], or diagnostic device require good initial guesses without operator intervention. We propose a method of mapping transient response data for a target system to a high-quality initial guess, or pre-estimate. The term pre-estimate emphasizes that the result is a preliminary or initial guess for a subsequent iterative minimization procedure.

Better initial guesses are not the only way to improve unsupervised estimator performance. Custom estimators incorporating insights from the model are often effective. Examples of specialized estimators for induction machine diagnostics are given in [2]. More recently, [3] gives techniques specialized to the multiple sinewave fitting problem. In [4], a special estimator is consider for the sinewave fitting problem. The drawback of custom schemes is that they require significant design effort. Additional work is required for testing, as in [5]. Genetic algorithms may also be used to search for good parameter values [6]. Genetic algorithms appear to offer a tradeoff where less knowledge of the system can be accommodated, but significant computational effort may be needed to find a satisfactory solution [6]. In contrast to the custom estimator approach, our pre-estimation scheme works with any identification method that needs a high-quality initial guess. On-line pre-estimation computations are minimal and occur in fixed time, unlike genetic methods. Pre-estimator performance is often good enough that subsequent methods approach quadratic convergence.

Manuscript received May 1, 2002; revised January 17, 2004. This work was supported in part by the National Science Foundation, in part by Ford Motor Company, in part by the MSU VP for Research, and in part by the Grainger and Landsman Foundations.

The authors are with the Department of Electrical and Computer Engineering, Montana State University, Bozeman, MT 59717 USA.

Digital Object Identifier 10.1109/TIM.2004.827299

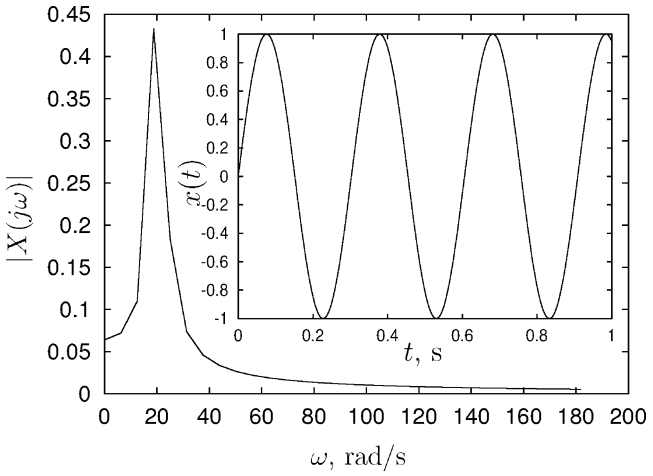

Fig. 1. Scaled FFT for sinewave frequency initial guess.

Experienced users usually pick an initial guess for interactive system identification based on rough estimates from the data. For example, consider identifying the frequency of sinewave from time-domain data as shown in Fig. 1. The pure sinewave is a simple example, but characteristic of more difficult problems that can be found in [3], [7]. A user would likely examine the data to find a good initial guess before starting an iterative system identification procedure. Zero crossings can provide a quick estimate of the frequency. Alternatively, the fast Fourier transform (FFT) shown in Fig. 1 might be used. Fig. 1 is scaled so that the amplitudes and frequencies correspond to the continuous-time Fourier transform of the signal; however, the spectrum is spread over a range of frequencies due to end effects (see [8] for better methods of estimating continuous-time spectral content using the FFT). Still, the FFT has a peak close to the system frequency, and could provide an excellent initial guess for a subsequent minimization procedure. Both the zero crossing and the FFT approach transform operator knowledge about the structure of the response (a sinusoid) directly to initial guess information.

Other strategies of computing system parameters directly from system responses have been proposed. Like the FFT applied to the sinewave problem, these could be used to form pre-estimates from observed data. One such approach is suggested in [9] for a nonlinear induction machine system identification problem. In [9], easy-to-identify, reduced-order low-time and high-time models are extracted from the full nonlinear induction machine model and used to obtain pre-estimates of the system parameters. This is analogous to the approach that might be used by a human operator familiar

with induction machine dynamics. Finding initial parameters without iteration is suggested by Caudill in [10] in terms of "directly inverting" a physical system. Caudill proposes finding the parameter vector estimate $\hat{\mu}$ from the inputs u and outputs y using a function

$$\hat{\mu} = F_S(u, y). \tag{1}$$

The existence of this function over some range of useful input and output signals is equivalent to the identifiability of the system given those inputs and outputs. Pre-estimates are intended to accelerate and assure convergence for well-formulated problems, so we assume the existence of F_S and seek an approximation to use for initial guesses. In [10], F_S was found explicitly for a particular system.

The quality required of a pre-estimate depends on the minimization procedure selected for the identification problem. For nonlinear least squares, methods that use second-order information about the Hessian may perform better than general-purpose Gauss–Newton (GN) or Levenburg–Marquardt (LM) methods in the large-residual, poor initial guess scenario [11]. However, like GN or LM, any optimization method that uses gradient or other localized information about the loss function can get stuck in local minima. Local minima are usually associated with undesired or unanticipated parameter estimates that can be especially troubling in unsupervised applications, for example, in embedded instrumentation. Pre-estimate quality should be assessed relative to a particular method, for example, by comparison of the number of iterations required with and without a pre-estimate.

Notation can be confusing because of the temptation to use the same variable to refer to the parameter, its appearance in a model, a final estimate of the parameter, or a pre-estimate. We typically use μ to refer to a vector of parameters, with components $\mu_1, \mu_2, \ldots, \mu_n$. Without any modifying marks, μ is used as the argument to models, loss functions and, in the appropriate context, refers to the true values of the system parameters. The notation $\tilde{\mu}$ is a pre-estimate of the true parameters μ. Similarly, the notation $\hat{\mu}$ refers to a final estimate of μ resulting from the minimization of an appropriate error criterion.

II. Approach

For many models, an analytical approach to finding a "model inverse" F_S is impractical. However, as emphasized by Fig. 2, a simulation model of the system establishes a relationship between the input $u(t)$, output $y(t)$, and parameters μ. Given a range of inputs and a range of parameter vectors μ, the simulation model can be used to explore the relationship between $u(t)$, $y(t)$, and μ for purposes of building an approximation to F_S. A good approximation of F_S can be used to obtain a pre-estimate $\tilde{\mu}$ of the parameter vector μ given $u(t)$ and $y(t)$.

A difficulty with this approach is that the approximation of F_S must map a large space, samples of $u(t)$ and $y(t)$, to a relatively small space containing μ. To be useful in applications, F_S must also achieve this mapping under conditions of additive noise. The first requirement suggests that the model for F_S will have many degrees of freedom. The second requirement suggests that the validation of a model for F_S will be difficult, i.e.,

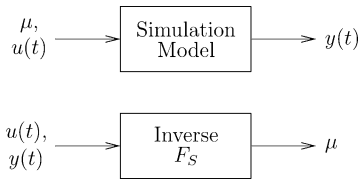

Fig. 2. Input–output relationships of the inverse model in comparison to the simulation model.

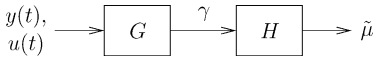

Fig. 3. Decomposition of inverse model into two steps. The first step G is a structured model that produces intermediate parameters γ with good noise properties. A good choice for G can reduce the degrees of freedom required in the black-box model H mapping γ to the desired parameters.

it is necessary to check that the expected behavior is preserved with noise. One possibility is to make F_S two parts, as suggested in Fig. 3. In Fig. 3, G turns input/output information into intermediate values γ, which are subsequently mapped by H to the desired pre-estimate $\tilde{\mu}$. The principle advantage of the two-step approach in Fig. 3 is that the requirements of noise immunity and data reduction (G) can be separated from the arbitrary function approximation (H) needed to obtain a pre-estimate from γ.

The steps G and H proposed in Fig. 3 can be related to the FFT pre-estimation approach proposed for Fig. 1. The pre-processing and data reduction steps of G are the FFT calculation and truncation of the spectrum to the anticipated range of frequencies. The frequency response in Fig. 1 is therefore a plot of γ in Fig. 3. Finding a pre-estimate from γ in the tone-estimation case involves picking the frequency with the largest coefficient. This is H in Fig. 3.

The intermediate parameters γ obtained from G in Fig. 3 should ideally reduce the observational data and be asymptotically unbiased with respect to the disturbances in the measurements. To satisfy these requirements, G could consist of system identification models with familiar properties, e.g., autoregressive or Box–Jenkins models of appropriate order [8], or, perhaps, a simple projection of the observations on a suitable collection of vectors. The performance of G can be assessed independent of the final mapping to parameters. For example, the degree to which γ is independent of the disturbance and captures key features of the observed data can be determined by applying G and reconstructing the observations from γ. Perhaps most important, the engineer can add physically relevant structure to G without directly applying a physical model. This is sometimes called "semi-physical modeling" [12]. The second stage in Fig. 3, H, maps intermediate parameters to a pre-estimate $\tilde{\mu}$. Unless the problem is trivial, such as the frequency pre-estimate for Fig. 1, H will generally involve a black-box function approximation structure. The number of degrees of freedom in H can be limited if G achieves an appropriate reduction of the input data. As an example, the input to G might be two thousand samples of the response of a system with four parameters. If G reduces these samples to ten coefficients in γ, then H need only map from a ten-dimensional space to the four dimensional space of $\tilde{\mu}$. For any given model structure for H, the ten-to-four mapping would involve fewer degrees of freedom than the two thousand to four mapping needed without data reduction in G.

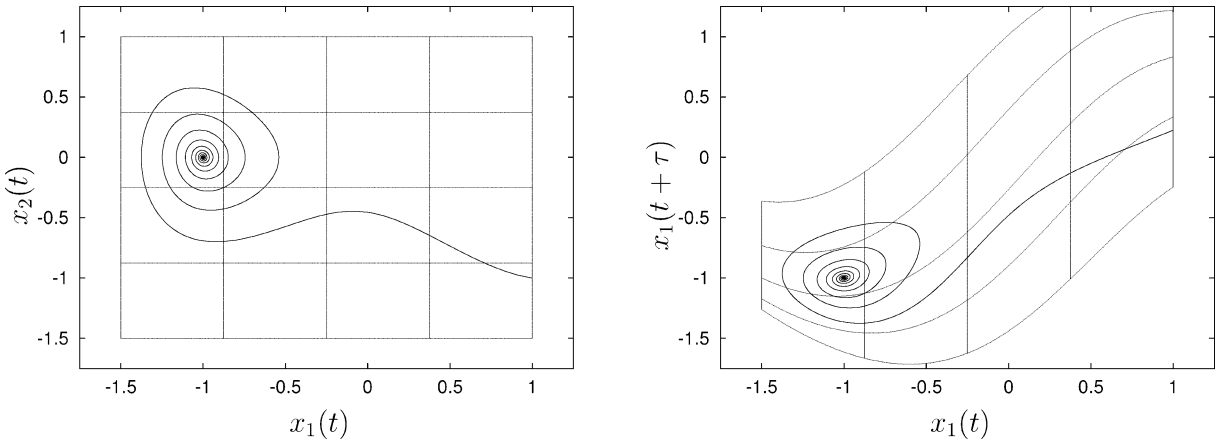

Fig. 4. State space trajectory $x_1(t)$, $x_2(t)$ of a Duffing oscillator compared to the lags of just one state variable. The state space and lag space are related by a smooth deformation, as illustrated by the grid lines.

Fewer degrees of freedom means that H can be found and tested with a realistic number of γ, μ pairs obtained by simulation.

While the two-step approach creates favorable conditions for approximating H, it does not suggest any particular model structure. If the observed data can be reconstructed from G and γ, and the system can be identified from the observed data, then in principle an H exists. This argument does not lead to a form for H. One approach is to try a generic function approximation model and validate its performance. Common choices that generalize cleanly to the multidimensional function approximation required of H include, for example, the two-layer perceptron network [13]–[16] and radial basis functions [12], [17]. Other structures, e.g., splines, might also be used. Global polynomials, historically a popular approximation tool, are probably a poor choice because their local approximation ability is obtained at the expense of rapidly diverging higher-order terms. These higher-order terms often lead to excessive interpolation errors with arbitrary data. See [17] or the preface to [18] for more details.

The following paragraphs offer possibilities for G and H that work well for problems we have explored. Design choices and perhaps even new structures will be needed for new problems because G and H are problem-specific.

A. Autoregressive Models for G

The idea of state-space reconstruction suggests a possibility for G in Fig. 3. State-space reconstruction (the "embedding theorem") holds that the state of a system (at least, the state that matters to the output) can be reconstructed from the lags of the inputs and outputs, up to a smooth local distortion [17]. The embedding theorem is especially interesting in its connection to the identifiability of a system. If a system is identifiable given data and a parameterization, by definition the parameters (and initial state if so parameterized) can be found from the input/output data. Given the initial state and parameters, the states can be reconstructed by simulation. In effect, the difference between plotting the lags of the response and identifying the parameters is finding the distortion that maps the lags to the states.

To illustrate the embedding theorem graphically, consider the Duffing oscillator [19]

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= x_1 - x_1^3 - \delta x_2. \end{aligned} \tag{2}$$

Fig. 4(a) shows a state-space trajectory $x_1(t)$, $x_2(t)$ for this system, with t ranging from 0 to 100 and $\delta = 0.2$. Fig. 4(b) shows the lagged pair $x_1(t)$, $x_1(t + 0.01)$. Fig. 4(b) is related by a mild distortion, or embedding, to the original states shown in Fig. 4(a). The rectangular grid in Fig. 4(a) makes the distortion of Fig. 4(b) clear.

State-space information from lagged data could greatly simplify parameter estimation problems. This is especially true in situations with a small number of sensors, where the observations are unlikely to include the entire state of the target system. Unfortunately, lagged data are not directly useful for finding state-space model parameters due to the unknown distortion between the lagged data and the state-space. However, lagged data can be combined with insight from the state-space model to design G. The relationship between the intermediate parameters produced by G and the desired system parameters can be embedded in H. Therefore, structures for G that include lagged data seem particularly promising.

The autoregressive (AR) model is a familiar structure incorporating lagged data linearly. The AR model has the form

$$y[k] + \gamma_1 y[k-1] + \gamma_2 y[k-2] + \cdots + \gamma_n y[k-n] = 0 \tag{3}$$

for some sampled signal $y[k]$ and coefficients γ_i. Given a record of y, this constraint can be immediately written as a linear system with rows

$$\begin{pmatrix} y[k-1] & y[k-2] & \cdots & y[k-n] \end{pmatrix} \begin{pmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_n \end{pmatrix} = -y[k] \tag{4}$$

which could be solved to find the coefficients γ_i. However, if y is a measurement, it will generally have some noise that will tend to bias estimates of the γ_i. Various methods exist for reducing the effect of this bias when estimating the AR parameters [8]. Although the AR model is linear, it may fit the responses of

certain nonlinear systems. An autoregressive G is used in the first example.

B. Radial Basis Functions for G and H

Radial basis functions are approximations using a superposition of functions depending on the distance r between the argument and a set of control points called centers. Radial basis functions may be appropriate for either H or G in Fig. 3. For a function $F(x)$, $F(R^P) \rightarrow R$, a scalar RBF approximation is

$$y = \sum_{i=1}^{N} a_i \phi\left(\|x - c_i\|_2\right) \tag{5}$$

for a set of N P-dimensional centers c_i [17]. Assuming that the same centers are appropriate for all components of a vector $F(x)$, or that the collection of centers is augmented until sufficient, a multidimensional RBF approximation is

$$y = A \begin{pmatrix} \phi\left(\|x - c_1\|_2\right) \\ \phi\left(\|x - c_2\|_2\right) \\ \vdots \\ \phi\left(\|x - c_N\|_2\right) \end{pmatrix} \tag{6}$$

where A is a matrix of coefficients and y is a column vector. Given the centers and a function $\phi(r)$, a coefficient matrix A can be determined by solving

$$\begin{pmatrix} \phi_{1,1} & \phi_{1,2} & \cdots & \phi_{1,N} \\ \phi_{2,1} & \phi_{2,2} & \cdots & \phi_{2,N} \\ \vdots & & \ddots & \vdots \\ \phi_{M,1} & \phi_{M,2} & \cdots & \phi_{M,N} \end{pmatrix} A' = \begin{pmatrix} F(x_1)' \\ F(x_2)' \\ \vdots \\ F(x_M)' \end{pmatrix} \tag{7}$$

where

$$\phi_{j,k} = \phi\left(\|x_j - c_k\|_2\right) \tag{8}$$

for M argument/output pairs x_j, $F(x_j)$. This method breaks down for $\phi(r)$ that introduce coefficients nonlinearly, e.g., $\phi(r) = e^{-r^2/a}$ with parameter a. However, for typical $\phi(r)$ like r, r^3, and e^{-r^2} [15], [17], determining the coefficients is a linear problem as suggested by (7). There are several strategies for picking the centers of a radial basis function [15], [17]. These include picking the centers by hand and relying on user intuition, parameterizing the centers to minimize some loss function over the training set (a nonlinear problem), picking arbitrary centers from the training data, picking uniform or random centers on or off the support of the data, or picking centers that are "representative" by clustering the training data.

Fig. 5 demonstrates several aspects of radial basis function approximation. In Fig. 5, the radial basis function approximates the scalar-valued time series $i_{qs}(t)$ corresponding to an induction motor startup transient. The argument of the approximation is t, and the centers c_i are indicated by the points. In this case, with $\phi(r) = r$, (7) becomes

$$\begin{pmatrix} |t_1 - c_1| & |t_1 - c_2| & \cdots & |t_1 - c_N| \\ |t_2 - c_1| & |t_2 - c_2| & \cdots & |t_2 - c_N| \\ \vdots & & \ddots & \vdots \\ |t_M - c_1| & |t_M - c_2| & \cdots & |t_M - c_N| \end{pmatrix} A' = \begin{pmatrix} i_{qs}(t_1) \\ i_{qs}(t_2) \\ \vdots \\ i_{qs}(t_M) \end{pmatrix} \tag{9}$$

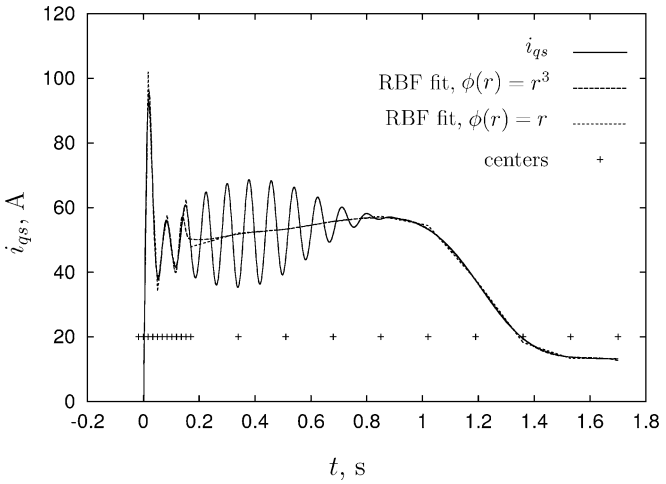

Fig. 5. Radial basis function approximation of induction motor transient response with linear and cubic $\phi(r)$. Centers for the radial basis function are indicated.

where A is a row vector. In Fig. 5, the approximation with $\phi(r) = r$ is a piecewise linear fit with the "joints" determined by the positions of the centers. With $\phi(r) = r^3$ and the same centers, the approximation is smoother. A final aspect of Fig. 5 is that the centers are selected for the qualitative goals of "good fit" in the detailed low-time portion of the transient, a "smoothed-over fit" in the middle section, and a close fit for the final part of the transient. In the case of the induction machine, these goals reflect that simplified high-time and low-time models of the transient can be used to deduce the full model parameters [9]. Important parts of the transient in Fig. 5 are accurately represented with a few coefficients. This makes the radial basis function attractive for G in Fig. 3. In general, the ease with which a user can control the radial basis function approximation by adjusting the centers is a useful property for G. Also, the simple procedure for finding the coefficients and straightforward generalization to higher dimensions make the radial basis function an attractive candidate for H in Fig. 3.

III. Simulation Results

A. Single-Tone Frequency Estimation

As an example, consider estimation of μ_1 given the model

$$y(t) = \mu_2 \sin(\mu_1 t) + \mu_3 \cos(\mu_1 t). \tag{10}$$

This signal is easily recognized as the response of a linear second-order differential operator, i.e.

$$\left[\frac{d^2}{dt^2} - \mu_1^2\right] y(t) = 0. \tag{11}$$

Therefore, a convenient and appropriate choice for G is a second-order discrete-time AR model

$$y[k-2] + \gamma_1 y[k-1] + \gamma_2 y[k] = 0 \tag{12}$$

where $y[k]$ is the sampled response $y[k] = y(kT)$ and T is the sampling interval. Generally, the relationship between μ and γ would be discovered by fitting H. However, in this case, the relationship between γ and μ can be found explicitly. If $y(t)$

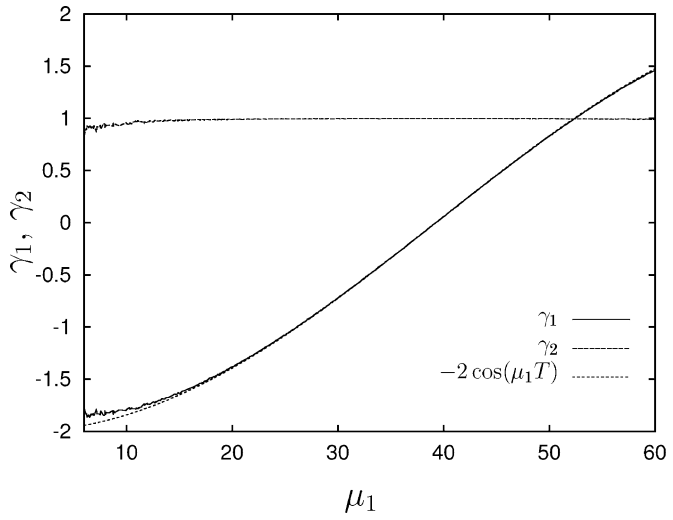

Fig. 6. Intermediate parameters γ_1, γ_2 from AR model as a function of μ_1.

satisfies (11), then the samples $y[k] = y(kT)$ also satisfy the discrete-time system

$$\left(z^{-2} - z^{-1} 2\cos(\mu_1 T) + 1\right) y[k] = 0 \tag{13}$$

where z^{-1} is the delay operator. By comparison of (13) with (12), the AR model coefficients are

$$\begin{aligned} \gamma_1 &= -2\cos(\mu_1 T) \\ \gamma_2 &= 1. \end{aligned} \tag{14}$$

If (14) were used to find μ_1 from γ_1, the appropriate branch of the $\cos^{-1}$ function would have to be selected in order to obtain μ_1 in the expected range. An approximation H to the inverse of (14) constructed using the desired values of μ_1 specifies the branch implicitly.

Applying the second-order AR model to a collection of 600 signals with $\mu_1 \in [6, 60]$, $T = .04$ s, each signal consisting of 500 samples, SNR of 10:1, and $\mu_2, \mu_3 \in [10, 16]$ produces the graph shown in Fig. 6. The theoretical $-2\cos(\mu_1 T)$ value for γ_1 is plotted in addition to the values obtained using the Matlab routine `ar` for γ_1 and γ_2. Note that, for smaller values of μ_1, the estimates for $\gamma_{1,2}$ are biased due to the presence of the normally distributed white noise. However, even with the bias, Fig. 6 suggests strongly that a curve-fitting technique might be used to extract pre-estimates of μ_1 from $\gamma_{1,2}$.

Four centers spaced evenly over the range of the γs were used in a RBF with $\phi(r) = r^3$ to map the intermediate parameters in Fig. 6 to the pre-estimate for μ_1. The RBF coefficients were obtained from the 600 signals used in Fig. 6. The performance of the AR/RBF pre-estimator is shown in Fig. 7. Fig. 7 shows the pre-estimate $\tilde{\mu}_1$ as a function of the actual value μ_1 for a set of 600 signals selected in a similar fashion but distinct from the signals used to obtain the RBF parameters. Fig. 7 also shows two lines indicating the range of values for which the output error least-squares loss function

$$V(\mu) = (y_{\text{obs}} - y(\mu_1))' (y_{\text{obs}} - y(\mu_1)) \tag{15}$$

is very nearly quadratic. The notation $y(\mu_1)$ is the simulated output of the system given the parameter μ_1, y_{obs} indicates the

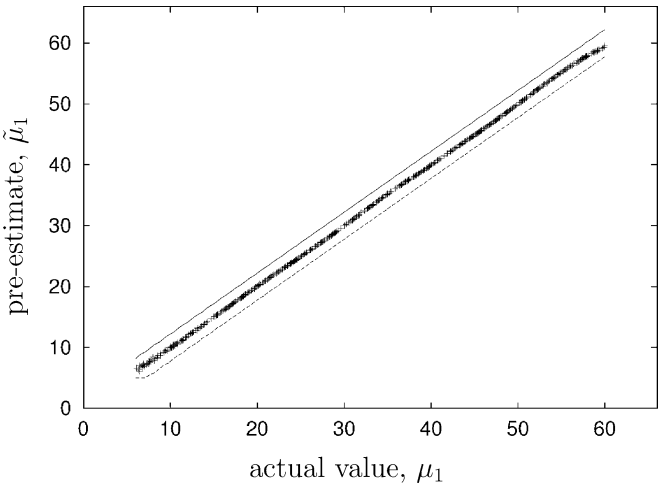

Fig. 7. Performance of AR/RBF pre-estimator for single-tone frequency estimation. Points show pre-estimate as a function of the true parameter. The lines above and below the points show the range of initial guesses for which a gradient method would quickly obtain the global minimum of the output-error, least-squares loss function.

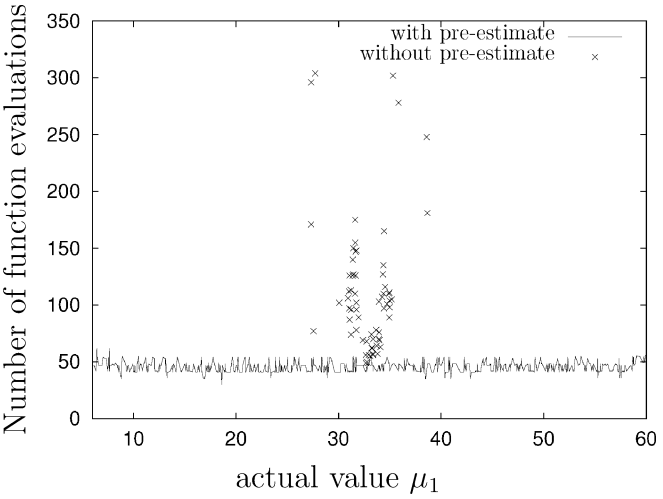

Fig. 8. Number of evaluations of the single-tone objective function for the Matlab routine `leastsq`. The line shows function evaluations needed with a pre-estimate, the points show evaluations required for a fixed initial guess.

observed values. As an approximation, the bounds in Fig. 7 assume correct values for μ_2 and μ_3 because these parameters can be determined linearly given μ_1. Fig. 7 suggests that typical nonlinear least squares routines would converge very quickly to the correct minimum for any one of the pre-estimates.

The most meaningful measure of pre-estimator performance is the degree to which the pre-estimates accelerate or ensure the convergence of an identification problem. Fig. 8 shows the number of evaluations of the objective function for parameter identification with and without a pre-estimate using the Matlab routine `leastsq`. The performance using pre-estimates is indicated by the solid line in Fig. 8 at about 50 function evaluations. An initial guess for `leastsq` consisting of the pre-estimate $\tilde{\mu}_1$ was used for the solid line, with $\tilde{\mu}_2$, $\tilde{\mu}_3$ determined by solution of the linear least-squares problem given $\tilde{\mu}_1$. The performance of `leastsq` without a pre-estimate is indicated in Fig. 8 by discrete data points. For these points, `leastsq` was started with the initial guess $(33 \quad 13 \quad 13)'$

corresponding to the average value of all the test parameter vectors. The parameters $\mu_{1...3}$ refer to the model in (10). With a pre-estimate (solid line in Fig. 8), about 50 function evaluations are required, the uncertainty in computation time is small, and the desired solution is obtained in all cases. Without a pre-estimate (points in Fig. 8), less than one hundred of the six hundred test problems converged within three hundred function evaluations. Three hundred is the default function evaluation limit imposed by `leastsq` for a three-parameter problem. In addition, Fig. 8 shows considerable variability in the number of function evaluations needed for the problems that did converge without a pre-estimate.

The results obtained without a pre-estimate in Fig. 8 can be understood in terms of the model and its loss function. Fig. 9 shows the least-squares loss function $V(\mu_1)$ defined by (15) and the model in (10), assuming *correct* values for μ_2 and μ_3 and a target solution of $\mu_1 = 20$. Only the narrow range of initial guesses for which the slope of the loss function leads to the minimum can be expected to converge to the desired answer using a typical identification method. In Fig. 9, an initial guess within about $\pm 5\%$ of the correct value is required. Similarly, in Fig. 9, the fixed initial guess with $\mu_1 = 33$ only converges to nearby problems.

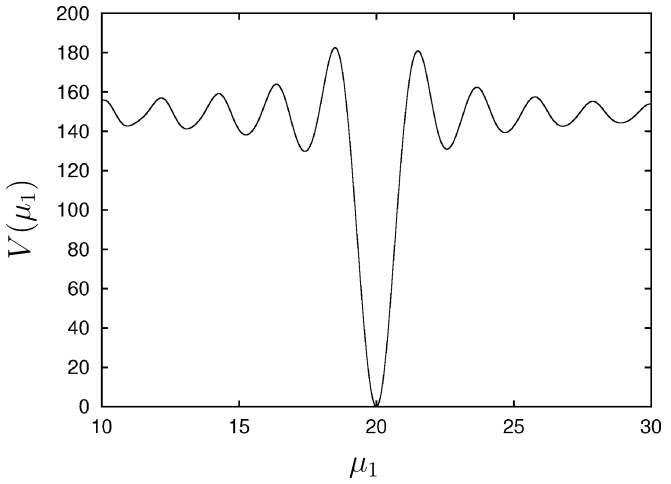

Fig. 9. Output error, least-squares loss function for convergence to $\mu_1 = 20$. The bounds in Fig. 7 correspond to the range between the peaks. Local minima show the need for a good initial guess even with this relatively simple problem.

B. Induction Machine Parameter Estimation

The dynamics of a three-phase induction machine can be modeled according to the synchronously rotating coordinate frame or dq-space equations

$$\frac{d}{dt}\begin{pmatrix}\lambda_{qs}\\ \lambda_{ds}\\ \lambda_{qr}\\ \lambda_{dr}\end{pmatrix} = \begin{pmatrix}v_{qs}\\ v_{ds}\\ 0\\ 0\end{pmatrix} - \begin{pmatrix}r_s i_{qs} + \omega_0 \lambda_{ds}\\ r_s i_{ds} - \omega_0 \lambda_{qs}\\ r_r i_{qr} + (\omega_0 - \omega_r)\lambda_{dr}\\ r_r i_{dr} - (\omega_0 - \omega_r)\lambda_{qr}\end{pmatrix} \quad (16)$$

where ω_0 is the frequency of excitation at the stator, ω_r is the rotor speed, and the λs are the flux linkages with rotor quantities and parameters as reflected to the stator [9], [20]. The voltages v_{qs} and v_{ds} are the excitation at the stator, and r_r and r_s are resistances associated with the rotor and stator, respectively. The flux linkages and currents are related according to

$$\begin{aligned}\lambda_{qs} &= L_l i_{qs} + (L_l + L_m)(i_{qs} + i_{qr})\\ \lambda_{ds} &= L_l i_{ds} + (L_l + L_m)(i_{ds} + i_{dr})\\ \lambda_{qr} &= L_l i_{qr} + (L_l + L_m)(i_{qs} + i_{qr})\\ \lambda_{dr} &= L_l i_{dr} + (L_l + L_m)(i_{ds} + i_{dr})\end{aligned}$$

where L_l is a leakage inductance and L_m is a magnetizing inductance. The mechanical part of the system determines how the speed and torque are related. For illustration assume a friction and mass mechanical model, i.e.

$$\frac{d}{dt}\omega_r = 3K(\tau - \beta\omega_r) \quad (17)$$

where $\tau = \lambda_{qr} i_{dr} - \lambda_{dr} i_{qr}$ is proportional to the torque of electrical origin and K is a parameter inversely proportional to the inertia connected to the motor shaft. The system parameter vector

$$\mu = (r_r \quad r_s \quad L_l \quad L_m \quad K \quad \beta)' \quad (18)$$

includes the damping β, inertia parameter K, and four electrical parameters r_r, r_s, L_l, and L_m. This model can be used to determine electrical parameters or for indirect identification of mechanical parameters using electrical measurements, as in [21].

A pre-estimation system was constructed using operations G and H as suggested in Fig. 3, both based on RBF approximations. A time-series RBF approximation with $\phi(r) = r^3$ was used for G, using the same centers as in Fig. 5 but applied to both i_{qs} and i_{ds}. The coefficients γ of this RBF approximation were output from G. For H, another RBF was used, with two centers for each coefficient output from G. These centers were selected to cover the range of γ generated when G was applied to 1500 induction motor simulations with randomly selected parameter vectors. The coefficients for H were obtained from 4000 randomly selected training parameter vectors and validated using another set of 1500 parameter vectors. The results are shown in Fig. 10, where pre-estimates for each parameter are plotted as a function of the true values.

For a single-parameter problem, it is convenient to show pre-estimate performance using upper and lower bounds corresponding to quadratic convergence as in Fig. 7. For a multidimensional problem, as in Fig. 10, the range of fast-converging values for a single parameter depends on the values of the other parameters. As a substitute, we find upper and lower scale factors $\alpha_{\min}$ and $\alpha_{\max}$ where any initial guess between $\alpha_{\min}\mu$ and $\alpha_{\max}\mu$ converges quadratically to μ for 90% of the parameter vectors in the validation set. These bounds give a rough indication of the sensitivity of the multidimensional estimation problem to the initial guess. The plots in Fig. 10 show bounds given by $\alpha_{\min}$ and $\alpha_{\max}$ as well as pre-estimate results as a function of the true parameter values.

The pre-estimates in Fig. 10 were also tested using the LM method [11], [22], [23]. LM converged in 61% of the 1500 validation test cases given a constant initial guess centered in the range of each parameter. With a constant initial guess, the Jacobian was evaluated 6884 times, requiring in excess of 41 000 induction motor transient simulations. Using pre-estimates as initial guesses, more than 99.9% of the 1500 test problems converged and the Jacobian was evaluated 4209

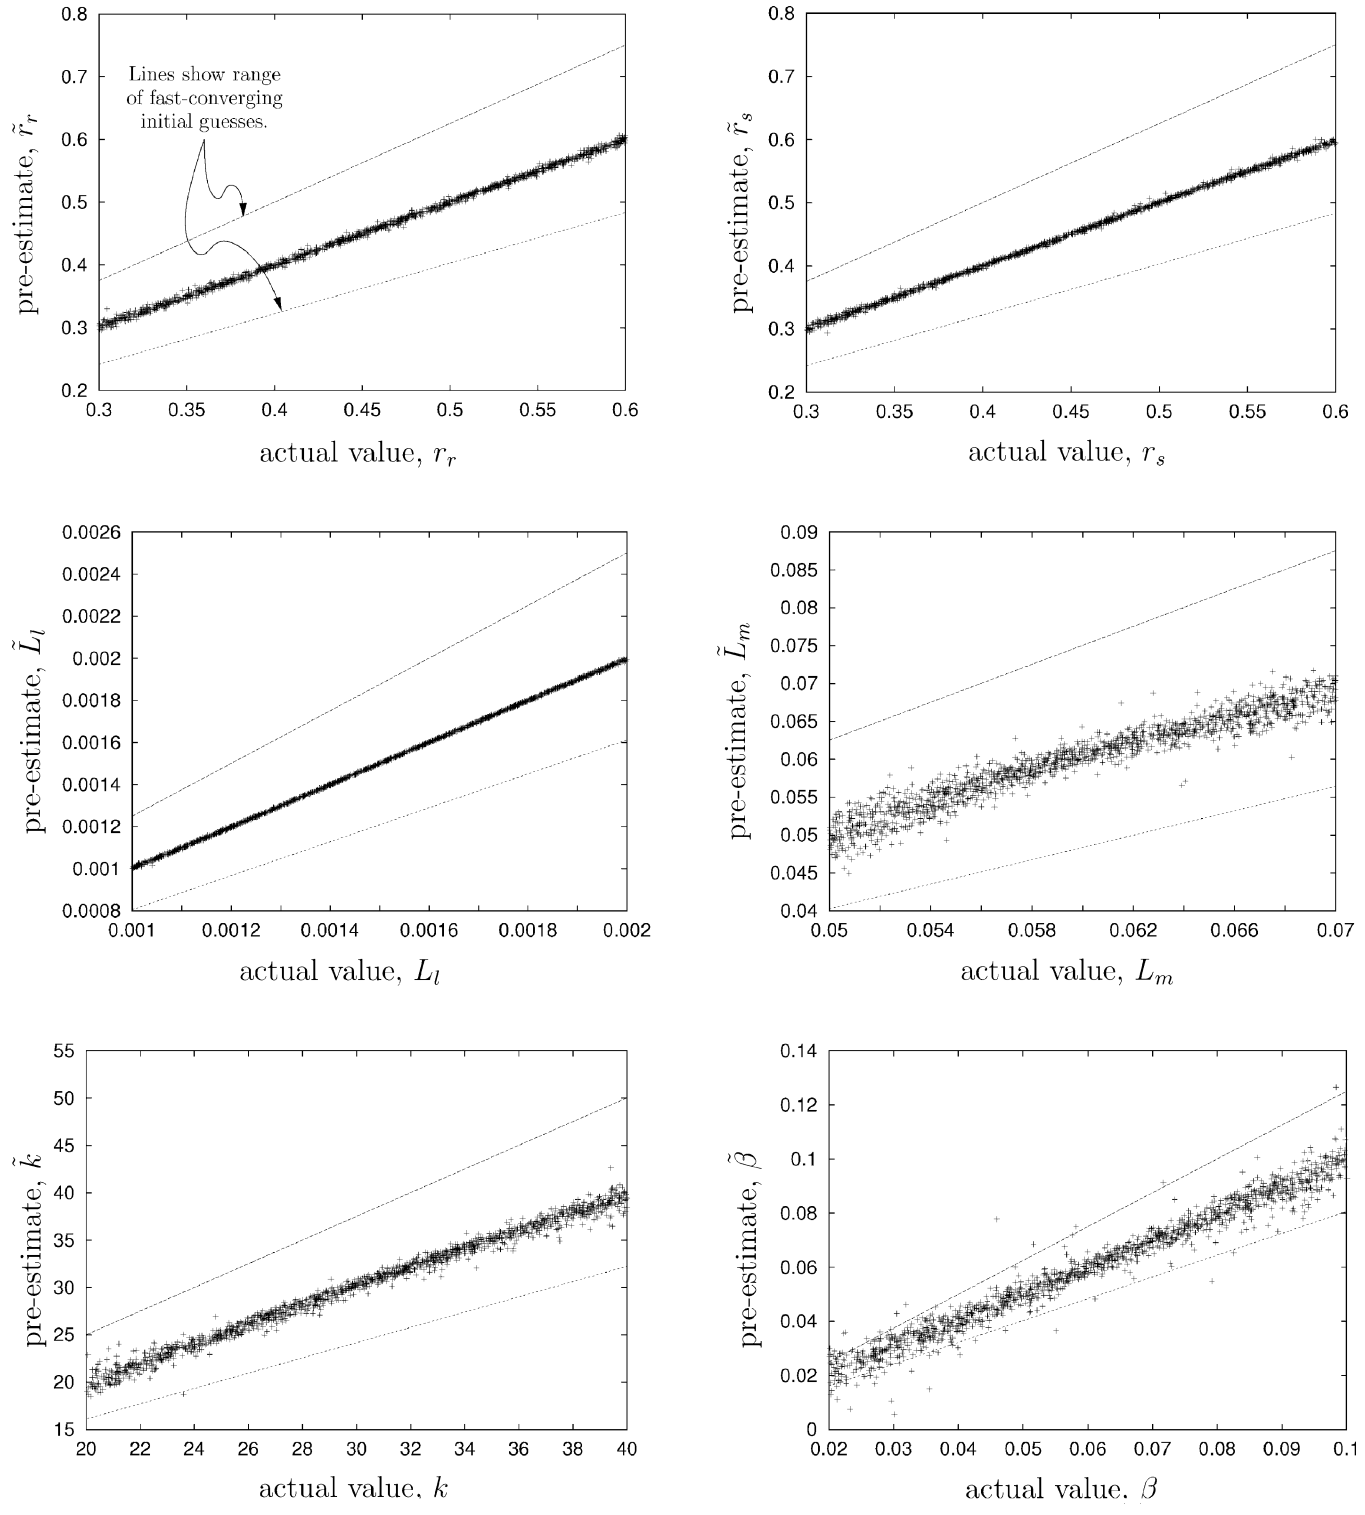

Fig. 10. Performance of RBF/RBF pre-estimator for induction machine models. Data points show performance with cross-validation data. The dashed lines indicate parameter ranges corresponding to fast convergence for subsequent minimization procedures.

times, requiring just over 25 000 transient simulations. This difference represents a significant performance improvement in many applications.

IV. Discussion

This paper explores a framework for computing high-quality initial guesses, or pre-estimates, from system responses. The pre-estimates demonstrated in this paper can be computed in fixed time and greatly accelerate subsequent iterative minimization. Pre-estimation is valuable in unattended or embedded applications because it expands the range of parameter values for which a conventional method will converge. In some cases, pre-estimation predictably reduces the number of iterations, as seen in Fig. 8.

Pre-estimation requires a certain degree of system knowledge. Like any conventional method, a system model is needed to build a relationship between the parameter space and the response. Rough ranges for the parameters must be provided. Excessively small parameter ranges will result in extrapolation errors, and excessively large parameter ranges will result in needlessly complex approximations. Finally, pre-estimation requires a user design process. The design challenge is to find appropriate fitting functions for G that reduce the data and have good noise performance. Pre-estimation fits in a spectrum of methods including genetic approaches and custom estimation routines that require varying degrees of system knowledge and design effort. Pre-estimation seems to offer a good tradeoff between performance and the need for prior information.

The two-step pre-estimation procedure suggested in this paper has fewer degrees of freedom than a one-step black-box map from system response to parameters. Still, there may be choices for G and H where interpolation errors cause spurious outputs. Similarly, extrapolation errors may result from system responses that fall outside of the design range. Errors of this kind are always a concern with black-box methods. However, errors in the pre-estimate are no worse than a bad initial guess. Pre-estimation may help ensure fast and predictable convergence in embedded applications where a good initial guess is not otherwise available.

ACKNOWLEDGMENT

The authors would like to thank Prof. Verghese, Prof. Kirtley, Prof. White, and Prof. Norford, for their valuable advice and support.

REFERENCES

[1] A. Ohsumi and N. Nakano, "Identification of physical parameters of a flexible structure from noisy measurement data," *IEEE Trans. Instrum. Meas.*, vol. 51, pp. 923–929, Oct. 2002.

[2] S. R. Shaw and S. B. Leeb, "Identification of induction motor parameters from transient stator current measurements," *IEEE Trans. Ind. Electron.*, vol. 46, pp. 139–149, Feb. 1999.

[3] G. Simon, R. Pintelon, L. Sujbert, and J. Schoukens, "An efficient nonlinear least square multisine fitting algorithm," *IEEE Trans. Instrum. Meas.*, vol. 51, pp. 750–755, Aug. 2002.

[4] J. Q. Zhang, Z. Xinmin, S. Jinwei, and H. Xiao, "Sinwave fit algorithm based on total least-squares method," in *Proc. IEEE Instrumentation and Measurement Technology Conf.*, Brussels, Belgium, June 1996, pp. 1436–1440.

[5] J. J. Moré, B. S. Garbow, and K. E. Hillstrom, "Testing unconstrained optimization software," *ACM Trans. Math. Software*, vol. 7, pp. 17–41, Mar. 1981.

[6] Y.-W. Chen, N. E. Mendoza, Z. Nakao, and T. Adachi, "Estimating wind speed in the lower atmosphere wind profiler based on a genetic algorithm," *IEEE Trans. Instrum. Meas.*, vol. 51, pp. 593–597, Aug. 2002.

[7] A. Routray, A. K. Pradhan, and K. P. Rao, "A novel kalman filter for frequency estimation of distorted signals in power systems," *IEEE Trans. Instrum. Meas.*, vol. 51, pp. 469–479, June 2002.

[8] R. Johansson, *System Modeling and Identification*. Englewood Cliffs, NJ: Prentice-Hall, 1993.

[9] S. R. Shaw, "Numerical methods for identification of induction motor parameters," M.S. thesis, Massachusetts Inst. Technol., Cambridge, Feb. 1997.

[10] L. F. C. Jr., H. Rabitz, and A. Askar, "A direct method for the inversion of physical systems," *Inverse Problems*, vol. 10, pp. 1099–1114, 1994.

[11] G. A. F. Seber and C. J. Wild, *Nonlinear Regression*. New York: Wiley, 1989.

[12] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1999.

[13] R. Williamson and U. Helmke, "Existance and uniqueness results for neural network approximations," *IEEE Trans. Neural Networks*, vol. 6, pp. 2–13, Jan. 1995.

[14] J. Sjöberg, Q. Zhang, L. Ljung, A. Benveniste, B. Delyon, P.-Y. Glorennec, H. kan Hjalmarsson, and A. Juditsky, "Nonlinear black-box modeling in system identification: a unified overview," *Automatica*, vol. 31, no. 12, pp. 1691–1724, 1995.

[15] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford, U.K.: Clarendon, 1995.

[16] T. Masters, *Practical Neural Network Recipes in C++*. New York: Academic, 1993.

[17] N. Gershenfeld, *The Nature of Mathematical Modeling*. Cambridge, U.K.: Cambridge Univ. Press, 1999.

[18] M. A. Kowalski, K. A. Sikorski, and F. Stenger, *Selected Topics in Approximation and Computation*. Oxford, U.K.: Oxford Univ. Press, 1995.

[19] S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. New York: Springer-Verlag, 1990.

[20] P. C. Krause, *Analysis of Electric Machinery*. New York: McGraw-Hill, 1986.

[21] S. R. Shaw, "System identification techniques and modeling for non-intrusive load diagnostics," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, Feb. 2000.

[22] J. J. Moré, B. S. Garbow, and K. E. Hillstrom, "User Guide for MINPACK-1," Argonne Nat. Lab., Argonne, IL, Tech. Rep. ANL-80-74, Aug. 1980.

[23] W. H. Press, S. A. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992.

Steven R. Shaw (M'00) received the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 2000.

He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Montana State University, Bozeman. He is interested in sensors, instrumentation, modeling, numerical and computational methods associated with control, and measurement problems.

Marc Keppler received the B.S. degree from Whitman College, Walla Walla, WA, in 2000. He is currently pursuing the M.S. degree in electrical engineering at Montana State University, Bozeman.

He is interested in control systems, instrumentation, and numerical methods for identification of physical systems.

Steven B. Leeb (SM'01) received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1993.

He has been a member of the MIT faculty in the Department of Electrical Engineering and Computer Science since 1993. He currently serves as an Associate Professor in the Laboratory for Electromagnetic and Electronic Systems, MIT. He is concerned with the design, analysis, development, and maintenance processes for all kinds of machinery with electrical actuators, sensors, or power electronic drives.