2018 Annual American Control Conference (ACC)
June 27–29, 2018. Wisconsin Center, Milwaukee, USA

Parameter Estimation for Nonlinearly Parameterized Gray-Box Models

Ankit Goel and Dennis S. Bernstein

***Abstract*— Many applications involve gray-box models, where the structure of the dynamics as a function of the parameters is known, but the values of the parameters are unknown. Nonlinear estimation algorithms, such as the extended Kalman filter (EKF) and the unscented Kalman filter (UKF), are typically applied to these problems. As an alternative approach, this paper uses retrospective cost model refinement (RCMR), which optimizes a retrospective cost function to update the gain of the estimator. In this paper, we investigate RCMR by estimating a single unknown parameter that may appear nonlinearly in linear and nonlinear systems.**

I. Introduction

Many applications involve gray-box models, where the structure of the dynamics as a function of the parameters is known, but the values of the parameters are unknown. These parameters are often *inaccessible*, which means that they cannot be measured directly due to sensor limitations. For example, the viscous force between an aerodynamic body and the contacting surface cannot be measured, and thus the drag coefficient is inaccessible. Consequently, measurements of the thrust and the motion of vehicle are needed along with a model in order to estimate the drag coefficient.

Nonlinear estimation algorithms, such as the extended Kalman filter (EKF) and the unscented Kalman filter (UKF), are typically applied to these problems. To do this, the uncertain parameters are modeled as constant states, which are appended to the original dynamics [1]–[4]. An alternative approach is to use the ensemble Kalman filter (EnKF) [5], [6], which can be viewed as a stochastic extension of UKF. Since the parameter states multiply the original states, the resulting dynamics are nonlinear whether or not the original dynamics are linear.

Yet another approach to parameter estimation is the variational method [7]–[9]. This approach requires an adjoint formulation of the dynamics and is computationally expensive since multiple iterations of the forward model and backward adjoint are required.

In [10], a two-step procedure for estimating the structured unknown parameters is suggested, where a black-box model is first constructed based on the input-output data, and a similarity transformation is used to recover the structured unknown parameters. This approach is limited to linear time-invariant systems.

The present paper focuses on retrospective cost model refinement (RCMR), which optimizes a retrospective cost function to update the gain of the estimator. Like UKF, RCMR uses the structure of the model with the current parameter estimates to propagate the states, but, unlike EKF (which requires the Jacobian of the dynamics), does not use the model for the parameter updates. Consequently, RCMR requires that the model update be computable, but the details of the computation need not be known. Also, unlike UKF and EnKF, RCMR does not require an ensemble of models. Finally, unlike adjoint-based methods, RCMR does not require an adjoint model.

RCMR can be used to estimate an inaccessible subsystem of an overall model, which includes uncertain parameters as a special case. This technique, developed in [11]–[13], uses an error signal given by the difference between the output of the physical system and the output of the model to update the parameter estimate. The parameter update is based on the retrospective cost function, whose minimizer updates the gain of the estimator. The ability to update a parameter or subsystem estimate provides a technique for data-based model refinement, which is also known as empirical model correction, model updating, and model calibration [14], [15].

In the present paper, we investigate the performance of RCMR in estimating parameters in a gray-box model. In particular, the contribution of the present paper is a numerical investigation of the performance of RCMR on a collection of examples involving a single unknown parameter where the functional parameter dependence is known and may be either linear or nonlinear.

The paper is structured as follows. In Section II, we formulate the problem of estimating unknown parameters in a gray-box model. In Section III, we present the retrospective cost parameter estimator structure. The RCMR algorithm is presented in Section IV. Numerical examples are presented in Section V. Finally, we conclude the paper with a discussion of the results and future work.

II. Problem formulation

Consider the multi-input, multi-output (MIMO) discrete-time physical system model

$$x(k+1) = f\left(x(k), u(k), \mu\right) + w_1(k), \tag{1}$$

$$y(k) = h\left(x(k), u(k), \mu\right) + w_2(k), \tag{2}$$

where $x \in \mathbb{R}^{l_x}$ is the state, $u \in \mathbb{R}^{l_u}$ is the input, $y \in \mathbb{R}^{l_y}$ is the measured output, $w_1 \in \mathbb{R}^{l_x}, w_2 \in \mathbb{R}^{l_y}$ are the process and measurement noise, and $\mu \in \mathbb{R}^{l_\mu}$ is the unknown parameter.

The functional forms of f and h are assumed to be known, and are either linear or nonlinear functions of the unknown

Ankit Goel is with the Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109, USA, ankgoel@umich.edu.

Dennis S. Bernstein is with the Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109, USA, dsbaero@umich.edu.

parameter μ. For example,

$$f(x,u,\mu) = \begin{bmatrix} \sin\mu & \cos\frac{\mu}{3} \\ \frac{e^{-\mu}}{3} & \frac{1}{1+\mu} \end{bmatrix} x + \begin{bmatrix} \log(1+\mu^2) \\ 1+\cos\mu \end{bmatrix} u, \tag{3}$$

$$h(x,u,\mu) = \begin{bmatrix} \mu & \mu^2 \end{bmatrix} x. \tag{4}$$

Next, we consider the estimation model

$$\hat{x}(k+1) = f\left(\hat{x}(k), u(k), \hat{\mu}(k)\right), \tag{5}$$

$$\hat{y}(k) = h\left(\hat{x}(k), u(k), \hat{\mu}(k)\right), \tag{6}$$

where $\hat{x}(k)$ is the computed state, $\hat{y}(k)$ is the computed output, and $\hat{\mu}(k)$ is the output of the parameter estimator at step k. The parameter estimator is updated by minimizing a cost function based on the performance variable

$$z(k) \triangleq \hat{y}(k) - y(k) \in \mathbb{R}^{l_y}. \tag{7}$$

The problem objective is to estimate μ using measurements of u and y. The parameter-estimation problem is represented by the block diagram in Figure 1.

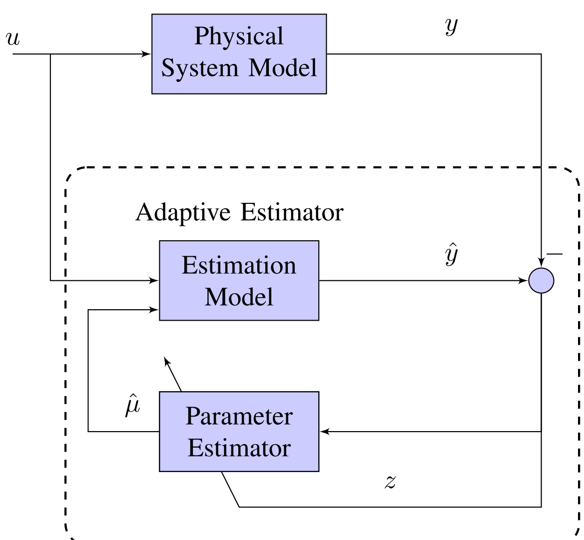

Fig. 1: Parameter-estimation architecture. The physical system modeled by the physical system model (1), (2) is driven by u and produces measurements y. The adaptive estimator consists of the estimation model (5), (6), which is driven by measurements of u and where the parameter estimate $\hat{\mu}$ is updated by the parameter estimator, which minimizes the error signal z.

III. Parameter Estimator

We consider a parameter estimator represented by an ARMA model with a built-in integrator. The parameter estimate $\hat{\mu}$ is thus given by

$$\hat{\mu}(k) = \sum_{i=1}^{n_c} P_i(k)\hat{\mu}(k-i) + \sum_{i=1}^{n_c} Q_i(k) z(k-i) + R(k) g(k), \tag{8}$$

where

$$g(k) = g(k-1) + z(k-1), \tag{9}$$

and $P_i(k) \in \mathbb{R}^{l_\mu \times l_\mu}$, $Q_i(k), R(k) \in \mathbb{R}^{l_\mu \times l_y}$ are the coefficient matrices, which are updated by the RCMR algorithm. The integrator is embedded in the estimator to ensure that $z(k) \to 0$ as $k \to \infty$ and thus, assuming identifiability and data persistency, that $\hat{\mu}(k) \to \mu$ as $k \to \infty$.

We rewrite (8) as

$$\hat{\mu}(k) = \Phi(k)\theta(k), \tag{10}$$

where the regressor matrix $\Phi(k)$ is defined by

$$\Phi(k) \triangleq I_{l_\mu} \otimes \phi^{\mathrm{T}}(k) \in \mathbb{R}^{l_\mu \times l_\theta},$$

where

$$\phi(k) \triangleq \begin{bmatrix} \hat{\mu}(k-1) \\ \vdots \\ \hat{\mu}(k-n_c) \\ z(k-1) \\ \vdots \\ z(k-n_c) \\ g(k) \end{bmatrix}, \tag{11}$$

$$\theta(k) \triangleq \mathrm{vec}\begin{bmatrix} P_1(k)\cdots P_{n_c}(k) & Q_1(k)\cdots Q_{n_c}(k) & R(k) \end{bmatrix} \in \mathbb{R}^{l_\theta},$$

$l_\theta \triangleq l_\mu^2 n_c + l_\mu l_y(n_c+1)$, "$\otimes$" is the Kronecker product, and "vec" is the column-stacking operator.

IV. RCMR Algorithm

In this section, we present the RCMR algorithm used to update the parameter estimator. RCMR is a specialized adaptation of the retrospective cost adaptive control (RCAC) algorithm [16].

A. Retrospective Performance Variable

We define the retrospective performance variable

$$\hat{z}(k) = z(k) + G_f(\mathbf{q})(\Phi(k)\hat{\theta} - \hat{\mu}(k)), \tag{12}$$

where $\mathbf{q}$ is the forward-shift operator, $\hat{\theta} \in \mathbb{R}^{l_\theta}$ contains the parameter estimator coefficients to be optimized,

$$G_f(\mathbf{q}) = \sum_{i=1}^{n_f} \frac{N_i}{\mathbf{q}^i}, \tag{13}$$

and, for all $i = 1, \ldots, n_f$, $N_i \in \mathbb{R}^{l_y \times l_\mu}$. G_f is an FIR filter of order n_f whose choice is discussed below. We rewrite (12) as

$$\hat{z}(k) = z(k) + N\Phi_b(k)\hat{\theta} - NU_b(k), \tag{14}$$

where

$$N \triangleq \begin{bmatrix} N_1 & \cdots & N_{n_f} \end{bmatrix} \in \mathbb{R}^{l_y \times n_f l_\mu},$$

$$\Phi_b(k) \triangleq \begin{bmatrix} \Phi(k-1) \\ \vdots \\ \Phi(k-n_f) \end{bmatrix} \in \mathbb{R}^{l_\mu n_f \times l_\theta},$$

$$U_b(k) \triangleq \begin{bmatrix} \hat{\mu}(k-1) \\ \vdots \\ \hat{\mu}(k-n_f) \end{bmatrix} \in \mathbb{R}^{l_\mu n_f}.$$

The vector $\hat{\theta}$, which contains the coefficients of the parameter estimator, is determined by minimizing the retrospective cost function, as described next.

B. Retrospective Cost Function

Using the retrospective performance variable $\hat{z}(k)$, we define the retrospective cost function

$$J(k,\hat{\theta}) \stackrel{\triangle}{=} \sum_{i=1}^{k} \hat{z}^{\mathrm{T}}(i) R_z \hat{z}(i) + \hat{\theta}^{\mathrm{T}} R_\theta \hat{\theta}, \tag{15}$$

where R_z and R_θ are positive definite. The following result uses recursive least squares (RLS) to minimize (15).

Proposition 4.1: Let $P(0) = R_\theta^{-1}$, $\theta(0) = 0$. Then, for all $k \geq 1$, the retrospective cost function (15) has a unique global minimizer $\theta(k)$, which is given by

$$\begin{aligned}\theta(k) = \theta(k-1) - P(k)\Phi_{\mathrm{b}}^{\mathrm{T}}(k)N^{\mathrm{T}} \\ \cdot \left(N\Phi_{\mathrm{b}}(k)\theta(k-1) + z(k) - NU_{\mathrm{b}}(k)\right),\end{aligned} \tag{16}$$

$$\begin{aligned}P(k) = P(k-1) - P(k-1)\Phi_{\mathrm{b}}^{\mathrm{T}}(k)N^{\mathrm{T}}\Gamma(k) \\ \cdot N\Phi_{\mathrm{b}}(k)P(k-1),\end{aligned} \tag{17}$$

where

$$\Gamma(k) \stackrel{\triangle}{=} \left(I_{l_z} + N\Phi_{\mathrm{b}}(k)P(k-1)\Phi_{\mathrm{b}}^{\mathrm{T}}(k)N^{\mathrm{T}}\right)^{-1}. \tag{18}$$

Furthermore, the parameter estimate is given by

$$\hat{\mu}(k) = \Phi(k)\theta(k). \tag{19}$$

C. The filter G_{f}

The cost function (15) can be written as

$$J(k,\hat{\theta}) = \hat{\theta}^{\mathrm{T}} A_\theta(k)\hat{\theta} + 2b_\theta(k)^{\mathrm{T}}\hat{\theta} + c_\theta(k), \tag{20}$$

where

$$A_\theta(k) \stackrel{\triangle}{=} \sum_{i=1}^{k} \Phi_{\mathrm{b}}^{\mathrm{T}}(i)N^{\mathrm{T}}N\Phi_{\mathrm{b}}(i) + R_\theta, \tag{21}$$

$$b_\theta(k) \stackrel{\triangle}{=} \sum_{i=1}^{k} \Phi_{\mathrm{b}}^{\mathrm{T}}(i)N^{\mathrm{T}}\left(z(i) - NU_{\mathrm{b}}(i)\right). \tag{22}$$

The batch least squares minimizer $\theta(k)$ of (15) is given by

$$\theta(k) = -A_\theta(k)^{-1}b_\theta(k). \tag{23}$$

The following result shows that the estimate $\hat{\mu}(k)$ of μ is constrained to lie in a subspace determined by the coefficients of G_{f}.

Theorem 4.1: Let $G_{\mathrm{f}}(\mathbf{q}) = \dfrac{N_1}{\mathbf{q}}$, $R_\theta = \beta I_{l_\theta}$, and $l_y = 1$. Then,

$$\hat{\mu}(k) = \alpha(k)N_1^{\mathrm{T}}, \tag{24}$$

where

$$\begin{aligned}\alpha(k) \stackrel{\triangle}{=} -\frac{1}{\beta}\sum_{i=1}^{k}\phi(k)^{\mathrm{T}}\phi(i-1)[z(i) - N_1\hat{\mu}(i-1) \\ + N_1\Phi(i-1)\theta(k)].\end{aligned} \tag{25}$$

Theorem 4.1 shows that the estimate $\hat{\mu}(k)$ produced by RCMR with a first-order FIR filter G_{f} is constrained to lie in the range of $N_1^{\mathrm{T}} \in \mathbb{R}^{l_\mu}$. If μ is scalar and $N_1 \neq 0$, then $\mathcal{R}(N_1^{\mathrm{T}}) = \mathbb{R}$. Therefore, we use a first-order FIR filter to estimate a scalar parameter μ. The numerical examples given in the next section show that the magnitudes of N_1 and β influence the convergence speed of the parameter estimate.

V. Numerical Examples

In this section RCMR is used to estimate an unknown scalar parameter that may appear nonlinearly in a system parameterization. Note that, like UKF, the functional form of the parameter dependence is used to propagate the parameter estimate; however, unlike EKF, explicit knowledge of the parameter dependence need not be known. In other words, the model must be computable, but the details of the computation need not be known by the user.

A. Linear Plant with Nonlinear Parameterization

Consider the LTI physical system model

$$x(k+1) = A(\mu)x(k) + B(\mu)u(k) + D_1w_1(k), \tag{26}$$

$$y(k) = C(\mu)x(k) + D_2w_2(k), \tag{27}$$

where $A(\mu), B(\mu), C(\mu)$ are given by (3), (4), the true value of μ is 0.8, $D_1 = \left[\begin{array}{cc} 0 & 1 \end{array}\right]^{\mathrm{T}}$, and $D_2 = 1$. The estimation model is

$$\hat{x}(k+1) = A(\hat{\mu}(k))\hat{x}(k) + B(\hat{\mu}(k))u(k), \tag{28}$$

$$\hat{y}(k) = E(\hat{\mu}(k))\hat{x}(k), \tag{29}$$

where $\hat{\mu}(k)$ is the output of the parameter estimator updated by RCMR.

We generate the measurement $y(k)$ using the input $u(k) = 2+\sin\left(\frac{2\pi}{40}k\right)+\sin\left(\frac{2\pi}{80}k - 0.3\right)+\sin\left(\frac{2\pi}{160}k - 0.5\right)$, the initial state $x(0) = [10\ 10]^{\mathrm{T}}$, the process noise $w_1 \sim \mathcal{N}(0, 10^{-4})$, and the measurement noise $w_2 \sim \mathcal{N}(0, 10^{-3})$. To reflect the absence of additional information, the initial state $\hat{x}(0)$ of the estimation model and the initial estimate $\hat{\mu}(0)$ of the unknown parameter μ are set to zero.

First, we use joint UKF as described in [3], [4] to estimate the states as well as the parameters of the physical system model (26), (27). We set the sigma point spread $\alpha = 0.1$, and the initial augmented state covariance $P_X(0) = I_3$. Figure 2 shows the estimates obtained using UKF. Figure 3 shows the parameter estimation error $\hat{\mu} - \mu$ at the end of the data window for various tuning choices in UKF. Note that the UKF estimates do not asymptotically converge to the true states and the unknown parameter. Furthermore, the accuracy of the estimates quickly degrades as the UKF tuning is varied.

Next, we use RCMR to estimate the unknown parameter μ in the linear system (26), (27) with the nonlinear parameter dependence (3), (4). We set $G_{\mathrm{f}}(\mathbf{q}) = \frac{1}{\mathbf{q}}$, $n_{\mathrm{c}} = 1$, $R_z = 1$, and $R_\theta = 10^8 I_{l_\theta}$. Figure 4 shows the estimate $\hat{\mu}(k)$ of μ.

Next, we investigate the effect of the choice of N_1 and R_θ on the parameter estimate. We set $n_{\mathrm{c}} = 1$, $R_z = 1$, and $R_\theta = 10^7 I_{l_\theta}$, and use RCMR to estimate μ with filter coefficient $N_1 \in \{10^{-1}, 10^{-0}, \ldots, 10^4\}$. Figure 5a shows the estimate $\hat{\mu}$ of μ for various filter coefficients. Next, we set $n_{\mathrm{c}} = 1$, $R_z = 1$, and $N_1 = 10^2$, and use RCMR to estimate μ with $R_\theta \in \{10^3 I_{l_\theta}, 10^5 I_{l_\theta}, \ldots, 10^{11} I_{l_\theta}\}$. Figure 5b shows the estimate $\hat{\mu}$ of μ for various parameter estimator coefficient weights. RCMR successfully estimates μ with filter coefficient N_1 and the estimator coefficient weight R_θ ranging over several orders of magnitude.

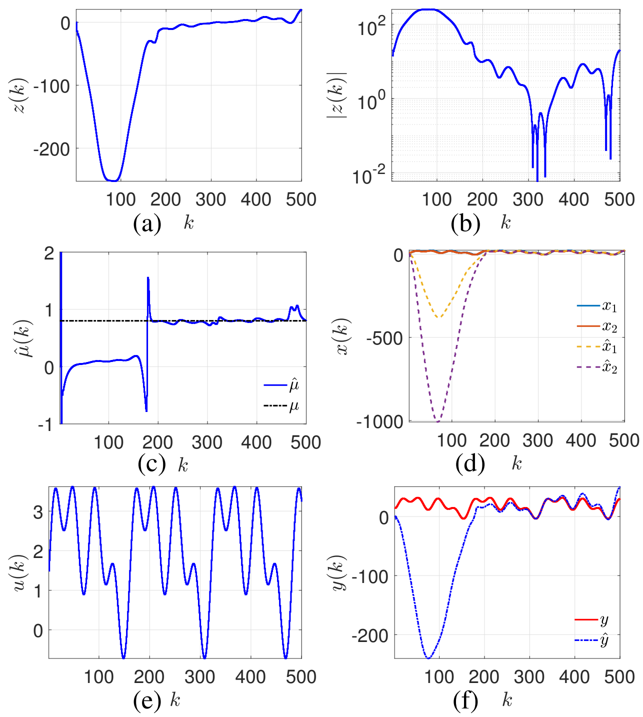

Fig. 2: UKF estimates of the states and the unknown parameter μ in the linear system (26), (27) with the nonlinear parameter dependence given by (3), (4). (a) shows the performance z on a linear scale, (b) shows the performance z on a log scale, (c) shows the parameter estimate $\hat{\mu}$, (d) shows the estimated states, (e) shows the measured input u to the system, and (f)

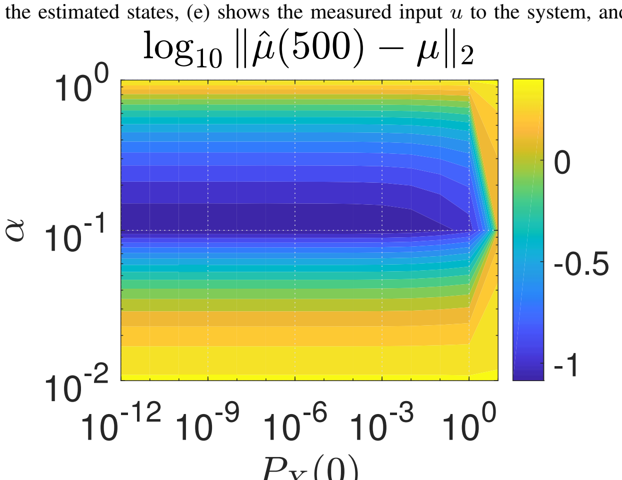

Fig. 3: Parameter estimation error $\hat{\mu} - \mu$ at the end of the data window for various values of the UKF tuning parameters.

Finally, we investigate the effect of the initial condition $x(0)$ of the physical system on the estimator performance. We generate the measurements y using the initial state $x_1(0) \in \{-100, -90, \ldots, 100\}$ and $x_2(0) \in \{-100, -90, \ldots, 100\}$, and set $n_c = 1$, $R_z = 1$, and $R_\theta = 10^8 I_{l_\theta}$, and use RCMR to estimate μ with filter coefficient $N_1 \in \{10^{-1}, 10^0, 10^1, 10^2\}$. Figure 6 shows the estimate error $\hat{\mu} - \mu$ at the end of the data window for various initial conditions of the physical system.

B. Damping of the Forced Van der Pol Oscillator

Next, we estimate the damping μ in the forced Van der Pol oscillator

$$\ddot{x} - \mu(1 - x^2)\dot{x} + x = u, \tag{30}$$

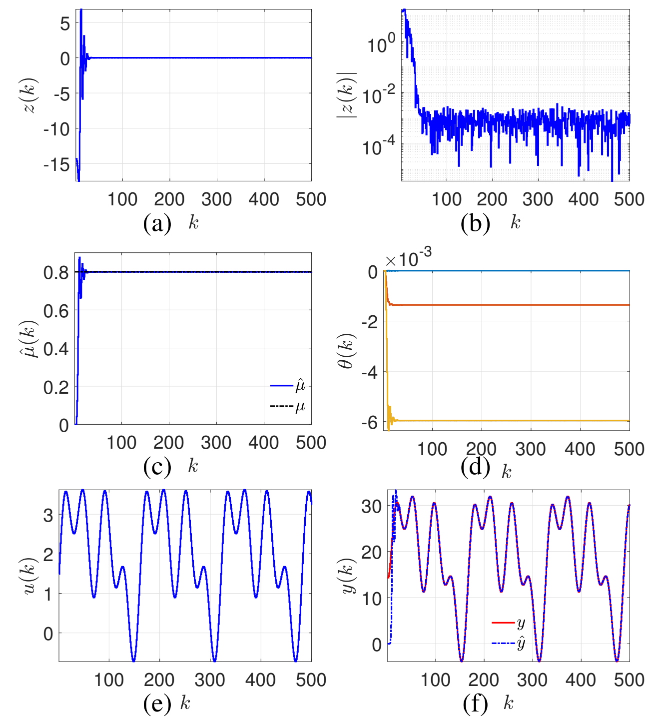

Fig. 4: RCMR estimate of the unknown parameter μ in the linear system (26), (27) with the nonlinear parameter dependence (3), (4). (a) shows the performance z on a linear scale, (b) shows the performance z on a log scale, (c) shows the parameter estimate $\hat{\mu}$, (d) shows the adapted coefficients θ of the parameter estimator, (e) shows the measured input u to the system, and (f) shows the measured output and the output of the estimation model.

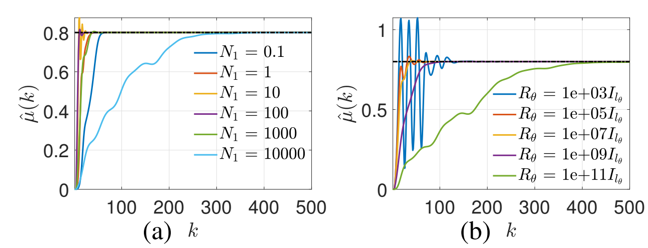

Fig. 5: Effect of weightings on the performance of RCMR. (a) shows the estimate $\hat{\mu}$ of μ for various filter coefficients, and (b) shows the estimate $\hat{\mu}$ of μ for various parameter estimator coefficient weights.

where u is the external forcing, and the unknown damping $\mu = 1.2$. We assume that x is measured, that is, $y = x$.

We write (30) in state-space form, and use Euler discretization with the time step $T_s = 1$ msec to generate discrete-time measurements $y(k)$ (denoting for convenience $y(.001k)$) with the input $u(k) = 20 \sin \frac{2\pi}{2000} k$, and the initial state $x(0) = [1\ 1]^{\mathrm{T}}$. The discrete-time equations are simulated and viewed as the truth model. The initial state $\hat{x}(0)$ of the model is set to zero, and the initial estimate of the damping $\hat{\mu}(0) = 0.5$. To prevent the RCMR estimate and the estimation model states from diverging, we project the parameter estimate $\hat{\mu}(k)$ onto $[0.5, 2]$.

We set $G_f(\mathbf{q}) = \frac{10^5}{\mathbf{q}}$, $n_c = 4$, $R_z = 1$, and $R_\theta = 10^5 I_{l_\theta}$. Figure 7 shows the estimate of the unknown damping μ. Figure 8 shows the actual and estimated state trajectory of the Van der Pol oscillator. Note that, although RCMR does not explicitly estimate the states of the system, the state estimates are driven to the actual states.

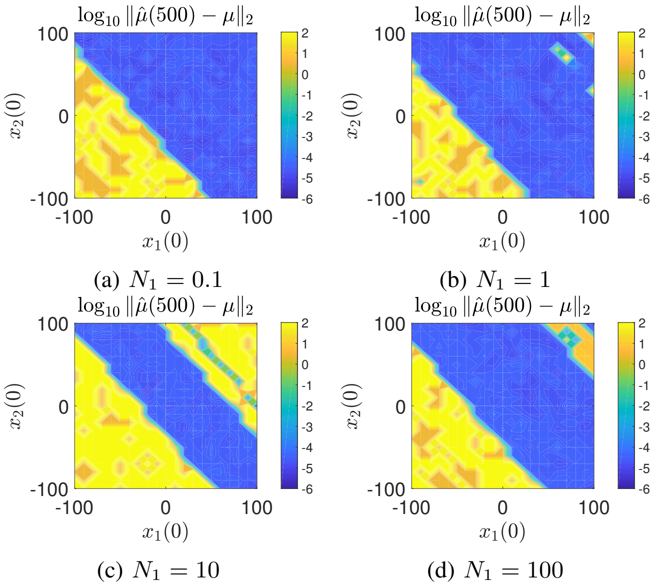

(a) $N_1 = 0.1$ (b) $N_1 = 1$

(c) $N_1 = 10$ (d) $N_1 = 100$

Fig. 6: Effect of the initial condition $x(0)$ of the physical system on the estimator performance. Theses plots show the estimation error $\hat{\mu} - \mu$ at the end of the data window for various choices of the filter coefficient N_1.

C. Translational Damping for a Pendulum on a Cart

We consider the problem of estimating the unknown damping coefficient μ in the translational dynamics of a pendulum on a cart modeled by

$$(m+M)\ddot{x} + \mu\dot{x} + mr\ddot{\phi}\cos\phi - mr\dot{\phi}^2\sin\phi = u, \quad (31)$$

$$r\ddot{\phi} + c\dot{\phi} + \ddot{x}\cos\phi + g\sin\phi = 0, \quad (32)$$

where x is the horizontal distance from the cart's center of mass to a point fixed in the ground, μ is the unknown translational damping coefficient, M is the mass of the cart, ϕ is the angle of the pendulum from the downward vertical direction, c is the damping coefficient of the pendulum, m is the mass of the pendulum, r is the length of the pendulum, and u is the horizontal force applied to the cart. The velocity of the cart is measured, that is, $y = \dot{x}$. Let $m = 10$ kg, $M = 1$ kg, $r = 1$ m, $\mu = 1$ N-sec/m, and $c = 1$ N-sec. We write (31), (32) in state-space form, and use Euler discretization with the time step $T_{\rm s} = 0.5$ msec to generate the measurements y with the input $u(k) = 0.5\cos\frac{2\pi k}{4000} + 0.1\cos\frac{2\pi k}{5000}$, and initial state set to $x(0) = \begin{bmatrix} 0.10 & -0.02 & 0.20 & -0.03 \end{bmatrix}$. The discrete-time equations are simulated and viewed as the truth model. The initial state $\hat{x}(0)$ of the model is set to zero, and the initial estimate of the damping is $\hat{\mu}(0) = 0$. To prevent the RCMR estimate and the estimation-model states from diverging, we project the parameter estimate $\hat{\mu}(k)$ onto $[-4, 4]$.

We set $G_{\rm f}(\mathbf{q}) = \frac{-0.1}{\mathbf{q}}$, $n_{\rm c} = 4$, $R_z = 1$, and $R_\theta = 10^{-1} I_{l_\theta}$. Figure 9 shows that the estimate $\hat{\mu}(k)$ converges to μ. However, it can be seen from Figure 10 that the state estimates do not converge to the actual states, which shows that RCMR does not correct the states, but rather corrects the unknown parameter so as to drive the output $\hat{y}$ of the estimation model to the measurements y.

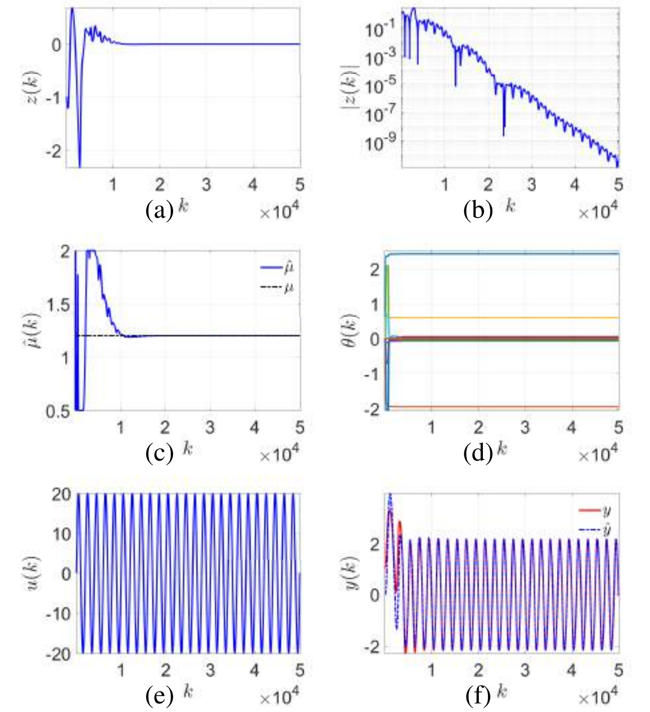

(a) (b) (c) (d) (e) (f)

Fig. 7: RCMR estimate of the unknown damping μ in the forced Van der Pol oscillator. (a) shows the performance z on a linear scale, (b) shows the performance z on a log scale, (c) shows the parameter estimate $\hat{\mu}$, (d) shows the components of the coefficient θ of the parameter estimator, (e) shows the measured input u of the system, and (f) shows the measured output and the output of the estimation model.

VI. Conclusions and Future Work

This paper numerically investigated the performance of RCMR for estimating an unknown parameter in a system with either linear or nonlinear parameter dependence. In particular, RCMR was used to estimate a parameter in a linear system whose realization has nonlinear parameter dependence, to estimate the damping in the Van der Pol oscillator, and to estimate the translational damping in a pendulum on a cart.

The analysis of RCMR focused on the structure of the filter $G_{\rm f}$ for which RCMR is potentially able to estimate the unknown parameter μ. In particular, if μ is scalar and $G_{\rm f}$ is a first-order FIR filter, then $\mathcal{R}(N_1^{\rm T}) = \mathbb{R}$. It thus follows from (24) that the subspace $\mathcal{R}(N_1^{\rm T})$, which contains $\hat{\mu}$, is independent of the choice of N_1. Consequently, in the case where μ is scalar and $G_{\rm f}$ is a first-order FIR filter, the choice of N_1 is not critical.

The situation is more complicated, however, in the case where μ is a vector. In this case, analysis of the retrospective cost showed that $G_{\rm f}$ must be higher order. Future work will thus focus on guidelines for constructing $G_{\rm f}$ for systems with two or more uncertain parameters. Additional objectives include integrating a state correction step to improve convergence, and investigating conditions on the input data to ensure consistency of the parameter estimator.

VII. Acknowledgments

This research was supported by AFOSR under DDDAS (Dynamic Data-Driven Applications Systems http://www.1dddas.org/) grant FA9550-16-1-0071.

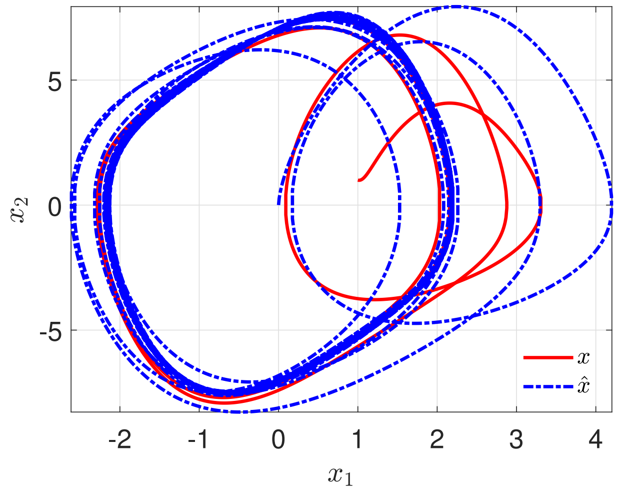

Fig. 8: Actual and estimated state trajectory of the Van der Pol oscillator. Note that RCMR does not seek to estimate the states of the system. However, this plot shows that the the state estimates converge to the actual states as a consequence of estimating the unknown parameter.

REFERENCES

[1] L. Ljung, "Asymptotic behavior of the extended kalman filter as a parameter estimator for linear systems," *IEEE Trans. Autom. Contr.*, vol. 24, pp. 36–50, 1979.

[2] G. L. Plett, "Extended kalman filtering for battery management systems of lipb-based hev battery packs: Part 3. state and parameter estimation," *J. Power Sources*, vol. 134, no. 2, pp. 277–292, 2004.

[3] R. Van Der Merwe and E. A. Wan, "The square-root unscented kalman filter for state and parameter-estimation," in *Int. Conf. Acoustics, Speech, Sig. Proc.*, 2001, pp. 3461–3464.

[4] E. A. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Adaptive Sys. Sig. Proc. Comm. Contr. Symp.*, 2000, pp. 153–158.

[5] G. Evensen, "The ensemble kalman filter for combined state and parameter estimation," *IEEE Contr. Sys. Mag.*, vol. 29, no. 3, 2009.

[6] H. Moradkhani, S. Sorooshian, H. V. Gupta, and P. R. Houser, "Dual state–parameter estimation of hydrological models using ensemble kalman filter," *Adv. Water Resources*, vol. 28, pp. 135–147, 2005.

[7] T. S. Jaakkola and M. I. Jordan, "Bayesian parameter estimation via variational methods," *Stat. Computing*, vol. 10, pp. 25–37, 2000.

[8] R. L. Raffard, K. Amonlirdviman, J. D. Axelrod, and C. J. Tomlin, "An adjoint-based parameter identification algorithm applied to planar cell polarity signaling," *IEEE Trans. Autom. Contr.*, vol. 53, pp. 109–121, 2008.

[9] M. Eknes and G. Evensen, "Parameter estimation solving a weak constraint variational formulation for an ekman model," *J. Geophys. Res.: Oceans*, vol. 102, pp. 12 479–12 491, 1997.

[10] G. Mercère, O. Prot, and J. A. Ramos, "Identification of parameterized gray-box state-space systems: From a black-box linear time-invariant representation to a structured one," *IEEE Trans. Autom. Contr.*, vol. 59, pp. 2873–2885, 2014.

[11] A. M. D'Amato, A. J. Ridley, and D. S. Bernstein, "Retrospective-cost-based adaptive model refinement for the ionosphere and thermosphere," *Stat. Anal. Data Mining*, pp. 446–458, 2011.

[12] X. Zhou, D. S. Bernstein, J. L. Stein, and T. Ersal, "Battery state of health monitoring by estimation of side reaction current density via retrospective-cost subsystem identification," *J. Dyn. Sys. Meas. Contr*, vol. 139, pp. 091 007–1–091 007–15, 2017.

[13] A. Goel, K. Duraisamy, and D. S. Bernstein, "Parameter estimation in the burgers equation using retrospective-cost model refinement," in *Proc. Amer. Contr. Conf.*, Boston, MA, 2016, pp. 6983–6988.

[14] J. Mottershead and M. Friswell, "Model updating in structural dynamics: a survey," *J. Sound Vibr.*, vol. 167, pp. 347–375, 1993.

[15] H. Madsen, "Automatic calibration of a conceptual rainfall–runoff model using multiple objectives," *J. Hydr.*, vol. 235, pp. 276–288, 2000.

[16] Y. Rahman, A. Xie, and D. S. Bernstein, "Retrospective Cost Adaptive Control: Pole Placement, Frequency Response, and Connections with LQG Control," *IEEE Contr. Sys. Mag.*, vol. 37, no. 5, pp. 28–69, October 2017.

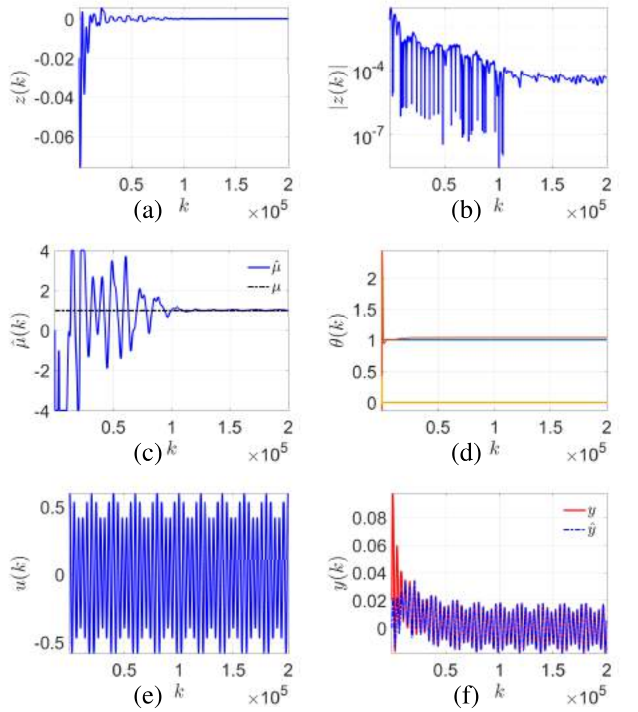

Fig. 9: RCMR estimate of the unknown translational damping μ of the cart. (a) shows the performance z on a linear scale, (b) shows the performance z on a log scale, (c) shows the parameter estimate $\hat{\mu}$, (d) shows the coefficients θ of the parameter estimator, (e) shows the measured input u to the system, and (f) shows the shows the measured output and the output of the estimation model.

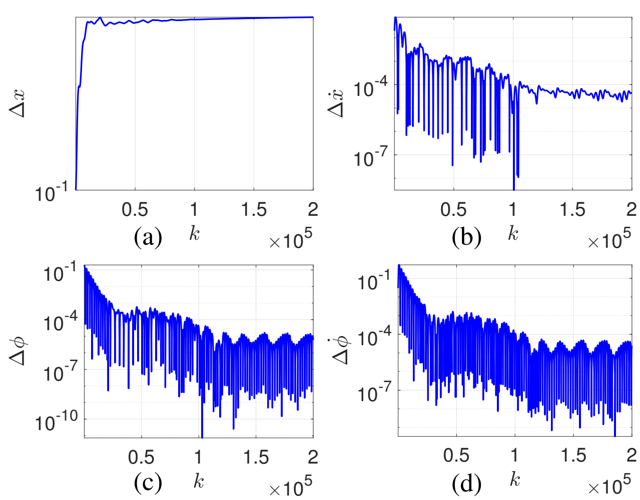

Fig. 10: State errors corresponding to the RCMR estimates of the unknown translational damping coefficient μ. Note that none of the state estimates converge to the actual states. Rather than estimating the states, RCMR corrects the unknown parameter in order to drive the output $\hat{y}$ of the estimation model to the measurements y.