

Neural Network Training Using Closed-Loop Data: Hazards and an Instrumental Variable (IVNN) Solution[★]

Johan Kon^{*} Marcel Heertjes^{*,**} Tom Oomen^{*,***}

^{*} *Control Systems Technology Group, Departement of Mechanical Engineering, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands, e-mail: j.j.kon@tue.nl*

^{**} *ASML, Mechatronics System Development, Veldhoven, The Netherlands*

^{***} *Delft Center for Systems and Control, Delft University of Technology, P.O. Box 5, 2600 AA Delft, The Netherlands*

Abstract: An increasing trend in the use of neural networks in control systems is being observed. The aim of this paper is to reveal that the straightforward application of learning neural network feedforward controllers with closed-loop data may introduce parameter inconsistency that degrades control performance, and to provide a solution. The proposed method employs instrumental variables to ensure consistent parameter estimates. A nonlinear system example reveals that the developed instrumental variable neural network (IVNN) approach asymptotically recovers the optimal solution, while pre-existing approaches are shown to lead to inconsistent estimates.

Keywords: Feedforward control, instrumental variables, neural networks.

1. INTRODUCTION

Improvements in feedforward control enable a major improvement in performance of control systems, e.g., in precision mechatronics. In feedforward control, two key requirements are typically considered (Clayton et al., 2009). First, high performance, i.e., a small tracking error, is desired. Second, task flexibility, i.e., performance for a variety of references, is required.

Learning techniques such as iterative learning control (ILC) (Moore, 1993) have allowed the generation of feedforward signals that achieve performance that can compensate all reproducible behaviour (Bristow et al., 2006; Oomen, 2020). Despite this high performance, these approaches lack task flexibility, as is evidenced by the development of ILC for flexible tasks, such as a signal library representing different subtasks (Hoelzle et al., 2011), and ILC with reference-dependent basis functions (van de Wijdeven and Bosgra, 2010; Boeren et al., 2018).

At the same time, traditional model-based feedforward control, which is highly flexible by design, has been further extended to enable higher performance. Advances include snap and static friction compensation through feedforward (Chang and Hori, 2003; Boerlage et al., 2004; Kontaras et al., 2017; van Haren et al., 2022) and rational feedforward to deal with higher-order flexibilities parametrizations (Bolder and Oomen, 2015; Boeren et al., 2018). These advances improve the performance and flexibility of the feedforward controller, but it is broadly experienced that these model-based extensions cannot achieve the same level of performance as learning-based techniques.

More flexible feedforward controller parametrizations have been investigated to go beyond the trade-off between performance and flexibility. For example, Gaussian processes

are employed as non-parametric feedforward controllers in Blanken and Oomen (2020); Poot et al. (2021). Additionally, Aarnoudse et al. (2021); Bolderman et al. (2021); Kon et al. (2022) use neural networks as parametric feedforward controllers. In Bolderman et al. (2021); Kon et al. (2022), neural networks are combined with a physics-based feedforward controller. In Aarnoudse et al. (2021), these neural networks are trained using learned ILC input signals for a variety of references, and closed-loop control aspects are taken into account through a control-relevant cost function. A key advantage of closed-loop or ILC data is that nonlinearities manifest themselves along the trajectory of interest (Markusson, 2001; Schoukens and Ljung, 2019), thus enabling to train a feedforward controller that is accurate in the domain of interest.

Although neural networks have been shown to enable major performance improvements for feedforward control of closed-loop systems, at present the complete role of noise in closed-loop situations has not been investigated. The aim of this paper is to illustrate that this noise may lead to inconsistent parameter estimates when training neural networks based on closed-loop data, and to provide a solution resulting in consistent estimates.

The results are illustrated for a specific class of systems and feedforward parametrizations for ease of exposition. The illustrated mechanism behind inconsistent estimates when using closed-loop data is independent of the specific neural network architecture, and holds for any neural network architecture, including convolutional neural networks and recurrent neural networks.

The main contribution of the present paper is an instrumental-variables (IV) approach to train neural networks based on closed-loop data in the presence of input disturbances. This generalizes the parametric results in feedforward tuning (Boeren et al., 2015, 2018) to the case where neural networks are used, and builds on parallel results in system identification (van Herpen et al., 2014) and ILC

[★] This work is supported by Topconsortia voor Kennis en Innovatie (TKI), and ASML and Philips Engineering Solutions.




Fig. 1. Standard feedforward control setup with input additive noise d (left) and feedforward control setup with neural network feedforward controller \mathcal{F}_φ (right).

(van Zundert et al., 2016). This contribution consists of the following cornerstones. First, the neural network feedforward parametrization is introduced in Section 2. Second, in Section 3, it is shown that estimation of the parametrization's coefficients with a least-squares cost function on closed-loop data results in estimates that do not approach a minimizer that result in the best performance for increasing data. Third, an instrumental-variable cost function is introduced in Section 4, for which the estimates do (locally) converge to the true coefficients. Lastly, in Section 5, these convergence properties and the consequences for performance are exemplified by simulation on a system with nonlinear friction characteristics.

Notation and Definitions

For the finite-time signal u with length N , $u(k) \in \mathbb{R}$ represents the signal at time index $k = 1, \dots, N$, and $\underline{u} = [u(1) \dots u(N)]^T \in \mathbb{R}^N$ the vector representation of this finite time signal.

2. PROBLEM FORMULATION

In this section, feedforward control for motion systems is introduced. Then, a neural network parametrization of the feedforward controller is given. The dataset to learn the coefficients of this parametrization is defined. Lastly, the problem of closed-loop noise in this dataset is formulated.

2.1 Feedforward Control for Motion Systems

The goal of feedforward control is to compensate for the effect of known exogenous inputs on the closed-loop system using an accurate feedforward signal. More specifically, consider the SISO control system shown in Fig. 1 (left). The goal of feedforward control is to determine feedforward signal $f(k) \in \mathbb{R}$ such that the output $y(k) \in \mathbb{R}$ of the discrete-time plant P equals the desired output $r(k) \in \mathbb{R}$ such that the error $e(k) \in \mathbb{R}$, given by

$$e(k) = r(k) - y(k) = r(k) - P(f(k)) \quad (1)$$

is zero $\forall k$, with $k \in \mathbb{Z}_{\geq 0}$ the time index. The feedback controller C aims to compensate the effects of both the unknown input disturbance d and plant dynamics that are not compensated by the feedforward signal.

If P is linear time-invariant (LTI), the error is given by

$$e = Sr + SP(f + d), \quad (2)$$

in which $S = (1 + CP)^{-1}$ is the sensitivity function, such that

$$f = P^{-1}r, \quad (3)$$

achieves perfect tracking of the reference, i.e., $e = 0$ in the noiseless ($d = 0$) setting and under appropriate initial conditions. Since P is not known exactly, various methods can be used to obtain f such as manual tuning, ILC with basis functions (Bolder and Oomen, 2015), or system identification (Söderström and Stoica, 1989) with inversion techniques (van Zundert and Oomen, 2018).




Fig. 2. Feedforward parametrization \mathcal{F}_φ as a combination of a two-lag delay line $D_2(q^{-1})$ and a multilayer perceptron with two fully connected layers.

However, P usually contains nonlinear dynamics. More specifically, P is an inverse nonlinear finite impulse response (NFIR) system, as is defined as follows.

Definition 1. The plant P with input $\tilde{u} \in \mathbb{R}$ satisfies the ordinary difference equation

$$P : \tilde{u}(k) \rightarrow y(k), \quad \tilde{u}(k) = g_y(D_m(q^{-1})y(k)), \quad (4)$$

with q^{-1} the forward shift operator, and

$$D_m(q)y(k) = [y(k) \ y(k-1) \ \dots \ y(k-m)]^T, \quad (5)$$

with $m \in \mathbb{Z}_{\geq 0}$ the maximum delay, and $g_y : \mathbb{R}^m \rightarrow \mathbb{R}$ an unknown, static, globally Lipschitz, nonlinear function.

Existing LTI feedforward parametrizations cannot capture all relevant dynamics of P , resulting in a loss of performance and necessitating a broader parametrization.

2.2 Neural Network Feedforward Parametrization

A neural network parametrization of the feedforward controller allows for learning all relevant plant dynamics from data, including unknown nonlinearities, since neural networks are universal approximators capable of approximating any nonlinear function. The feedforward controller \mathcal{F}_φ is parametrized by an NFIR system acting on the reference and its lags, see Fig. 1 (right) and 2, and is defined as follows.

Definition 2. The feedforward controller \mathcal{F}_φ satisfies

$$\mathcal{F}_\varphi : r(k) \rightarrow f(k), \quad f(k) = \mathcal{F}_\varphi(r(k)) = h^L(k), \quad (6)$$

in which h^L is the output of a neural network given by

$$\begin{aligned} h^l(k) &= D_p(q^{-1})r(k) && \text{if } l = 0 \\ h^l(k) &= \sigma(W^l z^{l-1}(k) + b^l) && \text{if } l = 1, \dots, L-1 \\ h^l(k) &= W^L z^{L-1}(k) + b^L && \text{if } l = L, \end{aligned} \quad (7)$$

with $W^l \in \mathbb{R}^{n_l \times n_{l-1}}$ the weights and $b^l \in \mathbb{R}^{n_l}$ the biases of layer l with n_l neurons, $\sigma(\cdot)$ an element-wise activation function, such as a sigmoid or hyperbolic tangent, and full parameter set $\varphi = \{W^l, b^l\}_{l=0}^L$.

The feedforward class \mathcal{F}_φ in (6) with $p = m$ can approximate g_y in (4) up to arbitrary accuracy, such that $|g_y(D_m(q^{-1})y(k)) - \mathcal{F}_\varphi(r(k))| < \epsilon \forall y(k) = r(k) \in \mathcal{R} \subset \mathbb{R}^m$ (Goodfellow et al., 2016). Given the correct parameters φ , \mathcal{F}_φ can generate high performance feedforward signals for a variety of references.

2.3 Closed-loop dataset

Estimation of the feedforward parameters in (6) requires a dataset describing the plant behaviour. A key advantage of closed-loop data over open-loop data is that the dataset captures the system dynamics around the trajectories of interest. Additionally, in many control applications, a

feedback controller is required to operate the system due to open-loop instability or safety requirements. Therefore, a closed-loop dataset \mathcal{D} covering the relevant IO space is considered, as formalized next.

Definition 3. The dataset $\mathcal{D} = \{\underline{u}, \underline{y}\}$, $\underline{u}, \underline{y} \in \mathbb{R}^N$ contains finite-time closed-loop measurements satisfying

$$y(k) = P(u(k) + d(k)) \quad u(k) = C(r(k) - y(k)). \quad (8)$$

A key disadvantage of a closed-loop dataset is that u and y contain spurious associations resulting from disturbance d , resulting in inconsistent estimates if not addressed appropriately.

Remark 4. Note that this closed-loop input data can also contain a contribution of a feedforward signal, obtained through, e.g., ILC. In fact, in the ILC case, the effect of disturbances is known to be amplified (Oomen, 2020, Section IV.B.4), amplifying the inconsistency problem.

2.4 Problem Formulation

The aim of this paper is to estimate φ of \mathcal{F}_φ in (6) based on the closed-loop dataset \mathcal{D} containing effects of input disturbance d such that, in the disturbance-free case, $P(\mathcal{F}_\varphi(r(k))) = r(k) \forall k \in \mathbb{Z}_{\geq 0}$. This includes

- 1) illustrating that existing least-squares (LS) criteria result in inconsistent parameter estimates (Section 3).
- 2) an IV criterion that results in consistent parameter estimates when initialized sufficiently close to the true parameters (Section 4), and
- 3) a simulation example illustrating that inconsistent LS estimates deteriorate performance (Section 5).

3. ANALYSIS OF LEAST-SQUARES CRITERION WITH CLOSED-LOOP DATA SUBJECT TO NOISE

This section illustrates that traditional approaches to neural network estimation can lead to poor estimation results. More specifically, it is shown that the inconsistency of the LS estimate also surfaces in the setting of neural network feedforward controllers estimated with closed-loop data. To this end, consider the LS criterion, which is commonly used for (nonlinear) regression in feedforward control with neural networks, defined as follows.

Definition 5. The least-squares criterion $J_{LS} \in \mathbb{R}_{\geq 0}$ is given by

$$J_{LS} = \sum_{k=1}^N (u(k) - \mathcal{F}_\varphi(y(k)))^2 = \|\underline{u} - \mathcal{F}_\varphi(\underline{y})\|_2^2, \quad (9)$$

in which $\mathcal{F}_\varphi(\underline{y}) \in \mathbb{R}^N$ is shorthand notation for

$$\mathcal{F}_\varphi(\underline{y}) = [\mathcal{F}_\varphi(y(1)) \ \mathcal{F}_\varphi(y(2)) \ \dots \ \mathcal{F}_\varphi(y(N))]^T. \quad (10)$$

For analysis purposes, it is assumed that feedforward parametrization (6) contains the inverse plant, i.e., $P^{-1} \in \mathcal{F}_\varphi$, as formalized by the following assumption.

Assumption 6. There exists a parameter set φ_0 such that

$$\mathcal{F}_{\varphi_0}(P(u)) = u \ \forall u \in \mathbb{R}. \quad (11)$$

This assumption is introduced merely to facilitate the forthcoming analysis. The main point of the analysis, i.e., the mechanism behind inconsistency for LS criterion (9), is present irrespective of this assumption. In practice, this assumption can only be asymptotically satisfied for $u \in \mathcal{U} \subset \mathbb{R}$ and $N_\varphi \rightarrow \infty$ by the universal approximation theorem (Goodfellow et al., 2016). Assumption 6 enables the following lemma.

Lemma 7. The least-squares criterion (9) is locally approximated around φ_0 by $J_{LS}^l \in \mathbb{R}_{\geq 0}$ given by

$$J_{LS}^l = \|\underline{u} - \mathcal{F}_{\varphi_0}(\underline{y}) - F(\underline{y})\Delta\varphi\|_2^2, \quad (12)$$

with $F(\underline{y}) = \frac{\partial}{\partial \varphi} \mathcal{F}_\varphi(\underline{y})|_{\varphi=\varphi_0}$ and $\Delta\varphi = \varphi - \varphi_0$.

The following assumption ensures that (12) has a unique minimum.

Assumption 8. $F^T(\underline{y})F(\underline{y})$ is nonsingular.

Assumption 8 imposes a persistence of excitation condition on y , such that φ can be uniquely determined. Lemma 7 and Assumption 8 allow for expressing the estimate $\hat{\varphi}_{LS}$.

Theorem 9. The estimate $\Delta\hat{\varphi}_{LS} = \arg \min_{\Delta\varphi} J_{LS}^l$ is given by

$$\Delta\hat{\varphi}_{LS} = \hat{\varphi}_{LS} - \varphi_0 = - (F^T(\underline{y})F(\underline{y}))^{-1} F^T(\underline{y})\underline{d}. \quad (13)$$

Next, the consistency of $\hat{\varphi}_{LS}$ is investigated, in which consistency is defined as follows.

Definition 10. A coefficient estimate $\hat{\varphi}$ of φ_0 based on N samples is consistent if (Söderström and Stoica, 1989), with probability 1,

$$\lim_{N \rightarrow \infty} \hat{\varphi} = \varphi_0. \quad (14)$$

A consistent estimate is both asymptotically unbiased and has diminishing variance for increasing data, i.e., converges to the true value with probability 1 for infinite data. Theorem 9 allows for analyzing the consistency of $\hat{\varphi}_{LS}$, as formalized in the following theorem.

Theorem 11. The least-squares estimate $\hat{\varphi}_{LS}$ is inconsistent, i.e.,

$$\lim_{N \rightarrow \infty} \hat{\varphi}_{LS} - \varphi_0 = \lim_{N \rightarrow \infty} \Delta\hat{\varphi}_{LS} \neq 0. \quad (15)$$

Theorem 11 shows that least-squares criterion (9) is inconsistent: even if the parametrization enables capturing the plant inverse, and $\hat{\varphi}_{LS}$ is initialized sufficiently close to φ_0 (or even at φ_0), φ_0 is not obtained as minimizer. Moreover, increasing noise levels increase the asymptotic bias in the parameter estimate.

4. INSTRUMENTAL VARIABLE NEURAL NETWORK

In this section, a new criterion for neural network training is presented that leads to consistent estimates by employing instrumental variables. It is shown that this IV criterion asymptotically recovers the true parameters if the optimization is initialized sufficiently close to the true parameters. To this end, consider the IV criterion defined as follows.

Definition 12. The instrumental-variables criterion $J_{IV} \in \mathbb{R}_{\geq 0}$ is given by

$$J_{IV} = \sum_{k=1}^N \|z(k)^T (u(k) - \mathcal{F}_\varphi(y(k)))\|_2^2, \quad (16)$$

with $z(k) \in \mathbb{R}^{1 \times N_\varphi}$, or equivalently

$$J_{IV} = \|Z^T (\underline{u} - \mathcal{F}_\varphi(\underline{y}))\|_2^2, \quad (17)$$

with $Z \in \mathbb{R}^{N \times N_\varphi}$ the matrix representation of $z(k)$.

The IV criterion is well-known in system identification and has demonstrated to be successful to deal with closed-loop issues (Söderström and Stoica, 1989) as well as in control (Boeren et al., 2018). The key novelty is its incorporation in neural network estimation, where LS criterion (9) is

standard practice. A key difference is the fact that neural networks are nonlinear models. To this end, a local analysis is performed in the remainder.

Lemma 13. The instrumental-variables criterion (16) can be locally approximated around φ_0 by $J_{IV}^l \in \mathbb{R}_{\geq 0}$ given by

$$J_{IV}^l = \|Z^T (\underline{u} - \mathcal{F}_{\varphi_0}(\underline{y}) - F(\underline{y})\Delta\varphi)\|_2^2. \quad (18)$$

The following assumption ensures that (18) has a unique minimum.

Assumption 14. $Z^T F(\underline{y})$ is nonsingular.

Assumption 14 implies that Z should be correlated with $F(\underline{y})$. Since $F(\underline{y})$ represents the parameter Jacobian of \mathcal{F}_φ , i.e., it is a nonlinear operation on \underline{y} , this is usually satisfied in practice by picking z such that it correlates with \underline{y} , as in the linear IV case (Söderström and Stoica, 1989). Degenerate cases, such as the case where $\mathcal{F}(\underline{y})$ is singular, could, in theory, occur, and require further analysis, but are considered outside the scope of this paper. Lemma 13 and 14 allow for expressing the estimate $\hat{\varphi}_{IV}$ of J_{IV}^l .

Theorem 15. The optimum $\Delta\hat{\varphi}_{IV} = \arg \min_{\Delta\varphi} J_{IV}^l$ is given by

$$\Delta\hat{\varphi}_{IV} = \hat{\varphi}_{IV} - \varphi_0 = -(Z^T(\underline{y})F(\underline{y}))^{-1} Z^T \underline{d}. \quad (19)$$

Theorem 15 allows for analyzing the consistency of $\hat{\varphi}_{IV}$ under the following additional assumption.

Assumption 16. Instruments z are uncorrelated with d .

The freedom that exists in the construction of Z readily allows for satisfying Assumption 16, e.g., through choosing Z as the plant output for a different noise realization, as the reference, or as a plant output predicted through a model. Under Assumption 16, consistency of $\hat{\varphi}_{LS}$ can be concluded, as formalized next.

Theorem 17. Given Assumptions 14 and 16, the estimate $\hat{\varphi}_{IV}$ is consistent, i.e., with probability 1,

$$\lim_{N \rightarrow \infty} \hat{\varphi}_{IV} - \varphi_0 = \lim_{N \rightarrow \infty} \Delta\hat{\varphi}_{IV} = 0. \quad (20)$$

Theorem 17 shows that instrumental-variables criterion (16) provides consistent estimates: if $\hat{\varphi}_{IV}$ is initialized sufficiently close to φ_0 , φ_0 is obtained as minimizer for infinite data.

5. SIMULATION EXAMPLE

In this section, the inconsistent parameter estimates of the LS cost function (9) are illustrated on an example dynamic system satisfying Definition 1. In contrast, it is shown that the IV criterion (16) produces consistent estimates when initialized sufficiently close to φ_0 . Furthermore, the consequences of the inconsistent estimate on the performance are illustrated.

5.1 Example System

The plant P is given by a mass-damper system with Stribeck-like friction characteristics, i.e., the kind of characteristics found in for example stage systems for lithographic inspection tools having a linear guidance with ball bearings, for which a simple model is given by

$$\tilde{u}(k) = m\delta^2 y(k) + c_1 \delta y(k) + \frac{c_2 - c_1}{\cosh(\alpha \delta y(k))} \delta y(k), \quad (21)$$

with parameters $c_1 = 1$, $c_2 = 20$, $\alpha = 2.5$, $m = 5$, and in which $\delta y(k) = T_s^{-1}(y(k) - y(k-1))$ represents the discrete-time derivative, with $T_s = 1/1000$. Note that $\delta^2 y(k)$ is a




Fig. 3. Stribeck-like friction curve (—) of example system (21) with $c_1 = 1$, $c_2 = 20$, $\alpha = 2.5$, consisting of a linear (---) and nonlinear contribution (····). This nonlinearity can be approximated up to arbitrary accuracy by \mathcal{F}_φ on the domain covered by r (---).




Fig. 4. Nominal reference (····), closed-loop output y_0 [m] (—) and scaled input u_0 [N] (---) for the noiseless setting.

linear combination of $y(k)$ and its $q = 2$ delayed instances $y(k-1)$ and $y(k-2)$, such that (21) satisfies Definition 1. The friction characteristics are visualized in Fig. 3.

The feedforward controller is parametrized as a neural network with 2 hidden layers with 10 neurons each and tanh activation functions, three input neurons and one linear output neuron, i.e.,

$$\mathcal{F}_\varphi(r(k)) = W_2 \sigma(W_1 \sigma(W_0 T D_2(q^{-1})r(k) + b_0) + b_1) + b_2, \quad (22)$$

with $\varphi = \{W_0, W_1, W_2, b_0, b_1, b_2\}$, $W_0 \in \mathbb{R}^{10 \times 3}$, $W_1 \in \mathbb{R}^{10 \times 10}$, $W_2 \in \mathbb{R}^{10 \times 1}$, $b_0, b_1 \in \mathbb{R}^{10}$, $b_2 \in \mathbb{R}$. $T \in \mathbb{R}^{3 \times 3}$ is a fixed matrix that transforms $D_2(q^{-1})r(k)$ into a derivative basis to alleviate training.

The system (21) is in feedback configuration, see Fig. 1, with controller $C(z)$ given by

$$C(z) = \frac{123.38z - 122.76}{z^2 - 1.908z + 0.91}. \quad (23)$$

Additionally, the system is subject to disturbance d with

$$d = H(z)\nu = \frac{0.8048z^2 - 1.61z + 0.8048}{z^2 - 1.57z + 0.65}\nu, \quad (24)$$

with ν white noise such that $\mathbb{E}\{\nu\} = 0$, $\mathbb{E}\{\nu^2\} = \sigma_\nu^2$.

The reference to be tracked is a fourth-order reference with $v_{max} = 0.5$, $a_{max} = 1$, $j_{max} = 62$, $s_{max} = 4100$. Fig. 4 shows the reference, the nominal noiseless output and the corresponding control effort (scaled).

To test the consistency of J_{LS} and J_{IV} , 20 noise realizations are generated for the standard deviation σ_ν ranging between 0 to 0.01. For each realization of the noise, a dataset consisting of the closed-loop responses u, y to the above reference is created. Each dataset is used to estimate




Fig. 5. Converged coefficient (3,1) of W_0 for $\hat{\varphi}_{LS}$ (\times) and $\hat{\varphi}_{IV}$ (\circ) for 20 noise realizations for a range of standard deviations σ_v , with means (—, —) visualized for convenience. Here, $\hat{\varphi}_{LS}$ does not correspond to the noiseless estimate (—), i.e., is biased, and has large variance, and thus is inconsistent. In contrast, $\hat{\varphi}_{IV}$ is consistent: $\hat{\varphi}_{IV}$ corresponds to the noiseless estimate, i.e., does not contain bias, and has small variance.

the parameters φ of the feedforward parametrization (22), both according to LS criterion (9) and IV criterion (16). For the IV criterion, the instrumental variables are given by the reference and its N_φ lags, i.e.,

$$z(k) = [r(k) \ r(k-1) \ \dots \ r(k-N_\varphi+1)]. \quad (25)$$

The parameters of each network are initialized at the parameters φ_0 of a network trained until convergence in the noiseless setting, i.e., at a perfect inverse up to approximation capabilities.

5.2 Parameter Inconsistency

Based on above datasets, the (in)consistency of $\hat{\varphi}_{LS}$ and $\hat{\varphi}_{IV}$ is demonstrated. Fig. 5 shows the converged parameter value of entry (3,1) of W_0 for the criterion J_{LS} and J_{IV} for each noise realization over a range of the noise standard deviation σ_v .

From Fig. 5, the following is observed.

- The estimate $\hat{\varphi}_{LS}$ does not correspond to the true parameter φ_0 of the noiseless setting for nonzero noise, since it is inconsistent. In contrast, $\hat{\varphi}_{IV}$ is a consistent estimate and, thus, for small nonzero noise levels, does converge to φ_0 up to minor (invisible) variations caused by the finite sample size N .
- For $\sigma_v = 0$, both $\hat{\varphi}_{LS}$ and $\hat{\varphi}_{IV}$ converge to φ_0 .

This inconsistency of $\hat{\varphi}_{LS}$, i.e., deviation from the noiseless parameters, indicates that traditional criteria may lead to an inaccurate inverse in the presence of noise due to systematic errors, as will be shown next.

5.3 Performance loss

The inconsistency of $\hat{\varphi}_{LS}$ also results in a loss of performance. Fig. 6 shows the norm of residuals between the optimal input f_0 and the generated input $\mathcal{F}_\varphi(y_0)$, in which f_0 and y_0 represent the plant input and output in the noiseless case to avoid extrapolation errors.

From Fig. 6, the following is observed.

- For zero noise, $\hat{\varphi}_{LS}$ and $\hat{\varphi}_{IV}$ result in the same performance. This is to be expected, since the estimates converge to almost exactly the same value (see Fig. 5) due to the low noise level.
- For $\sigma_v \in (0, 0.008]$, $\hat{\varphi}_{LS}$ suffers from reduced performance and more variation in performance due to




Fig. 6. Norm of residuals between true input f_0 and generated input $\mathcal{F}_\varphi(y_0)$ for $\hat{\varphi}_{LS}$ (\times) and $\hat{\varphi}_{IV}$ (\circ) with noiseless output y_0 as test reference, for 20 noise realizations for a range of standard deviations σ_v . The inconsistency of $\hat{\varphi}_{LS}$ causes a deviation of $\mathcal{F}_{\hat{\varphi}_{LS}}(y_0)$ from f_0 . $\hat{\varphi}_{IV}$ is locally consistent, and for small enough level, performs as well as the noiseless estimate (—).




Fig. 7. Normalized residuals between true input f_0 and generated input $\mathcal{F}_{\hat{\varphi}_{LS}}(y_0)$ (—) and $\mathcal{F}_{\hat{\varphi}_{IV}}(y_0)$ (—), corresponding to the realization of median performance for noise level $\sigma_v = 0.005$. $\mathcal{F}_{\hat{\varphi}_{LS}}(y_0)$ has larger deviations from the optimal input f_0 than $\mathcal{F}_{\hat{\varphi}_{IV}}(y_0)$ due to the inconsistent estimates $\hat{\varphi}_{LS}$.

the inconsistent estimates that depend on the noise realization. Additionally, higher noise levels result in more performance deterioration.

- For $\sigma_v > 0.002$, other effects besides (in)consistency of the estimator, discussed at the end of this section, start playing a role, such that the instrumental variable estimate is only locally able to obtain better estimates, and even performs worse for $\sigma_v > 0.008$.

Additionally, Fig. 7 shows the residuals between the optimal input signal f_0 and the generated feedforward signals $\mathcal{F}_{\hat{\varphi}_{LS}}(y_0)$ and $\mathcal{F}_{\hat{\varphi}_{IV}}(y_0)$ for the realization with median performance in terms of approximation norm of the optimal input for noise level $\sigma_v = 0.005$, see Fig. 6. It is observed that $\mathcal{F}_{\hat{\varphi}_{LS}}(y_0)$ is a worse approximation of f_0 than $\mathcal{F}_{\hat{\varphi}_{IV}}(y_0)$ caused by inconsistency of $\hat{\varphi}_{LS}$. These deviations from the optimal feedforward subsequently result in tracking errors.

Thus, the inconsistent estimate $\hat{\varphi}_{LS}$ results in performance deterioration, and may be improved through instrumental variables. When the noise level invalidates these local results, additional phenomena start playing a role.

- 1) Extrapolation errors occur, since the input during training covers a different part of g_y than the input during testing because the noise changes the trajectory of the closed-loop system. This requires that \mathcal{F}_φ

extrapolates, and the quality of this extrapolation depends on the parameters.

- 2) The quality of the local minima to which training converges, which depends on the noise realization, influences the performance.
- 3) Most importantly, the choice of instruments represents an implicit weighting, potentially not penalizing residuals in specific time intervals, which reflects in the quality of the generated feedforward signal in these segments. More specifically, when the instruments are chosen as the reference and its delayed values, time segments of the residual for which this reference is 0 are not penalized, e.g., $t \in [0, 0.5]$ s in Fig. 7. This allows for an offset in $\mathcal{F}_\varphi(y_0)$ in these intervals (e.g., through the bias parameter of the last layer of \mathcal{F}_φ), if only this offset results in smaller residuals for the penalized intervals. This is the mechanism behind the performance degradation of the IV criterion in Fig. 6.

These effects limit the applicability of the IV criterion outside a simulation context, and require a global analysis, which is part of future research.

6. CONCLUSION

In this paper, it is shown that the straightforward application of the least-squares criterion for learning neural network feedforward controllers from closed-loop data may introduce inconsistent parameter estimates. Importantly, these errors may go unnoticed in naively training neural networks. Furthermore, it is shown that this inconsistency can result in performance deterioration in terms of the generated feedforward signal. An instrumental variable approach for training neural networks results in consistent parameter estimates, overcoming previous shortcomings, and results in superior feedforward performance when initialized sufficiently close to the true optimum. These results are exemplified by simulation on a representative system with nonlinear friction characteristics. Future work focuses on globally consistent parameter estimation for closed-loop data, addressing local minima, extrapolation errors, and the choice of instrumental variables.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the detailed discussions with Leontine Aarnoudse.

REFERENCES

- Aarnoudse, L., Ohnishi, W., Poot, M., Tacx, P., Strijbosch, N., and Oomen, T. (2021). Control-Relevant Neural Networks for Intelligent Motion Feedforward. In *IEEE Int. Conf. Mechatronics*.
- Blanken, L. and Oomen, T. (2020). Kernel-based identification of non-causal systems with application to inverse model control. *Automatica*, 114(1).
- Boeren, F., Blanken, L., Bruijnen, D., and Oomen, T. (2018). Optimal Estimation of Rational Feedforward Control via Instrumental Variables: With Application to a Wafer Stage. *Asian J. Control*, 20(3), 975–992.
- Boeren, F., Oomen, T., and Steinbuch, M. (2015). Iterative motion feedforward tuning: A data-driven approach based on instrumental variable identification. *Control Eng. Pract.*, 37, 11–19.
- Boerlage, M., Tousain, R., and Steinbuch, M. (2004). Jerk derivative feedforward control for motion systems. *Proc. Am. Control Conf.*, 5, 4843–4848.
- Bolder, J. and Oomen, T. (2015). Rational basis functions in iterative learning control - With experimental verification on a motion system. *IEEE Trans. Control Syst. Technol.*, 23(2), 722–729.
- Bolderman, M., Lazar, M., and Butler, H. (2021). Physics-guided neural networks for inversion-based feedforward control applied to linear motors. *Conf. Control Technol. Appl.*, 1115–1120.
- Bristow, D.A., Tharayil, M., and Alleyne, A.G. (2006). A survey of iterative learning control. *IEEE Control Syst. Mag.*, 26(3), 96–114.
- Chang, B.H. and Hori, Y. (2003). Trajectory design considering derivative of jerk for head-positioning of disk drive system with mechanical vibration. In *Proc. Am. Control Conf.*, volume 5, 4335–4340.
- Clayton, G.M., Tien, S., Leang, K.K., Zou, Q., and Devasia, S. (2009). A Review of Feedforward Control Approaches in Nanopositioning for High-Speed SPM. *J. Dyn. Syst. Meas. Control*, 131(6).
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Hoelzle, D.J., Alleyne, A.G., and Wagoner Johnson, A.J. (2011). Basis Task Approach to Iterative Learning Control With Applications to Micro-Robotic Deposition. *IEEE Trans. Control Syst. Technol.*, 19(5), 1138–1148.
- Kon, J., Bruijnen, D., van de Wijdeven, J., Heertjes, M., and Oomen, T. (2022). Physics-Guided Neural Networks for Feedforward Control: An Orthogonal Projection-Based Approach. In *Proc. Am. Control Conf.*, To be published.
- Kontaras, N., Heertjes, M., Zwart, H., and Steinbuch, M. (2017). A compliance feedforward scheme for a class of LTV motion systems. In *Proc. Am. Control Conf.*, 4504–4509.
- Markusson, O. (2001). *Model and system inversion with applications in nonlinear system identification and control*. Ph.D. thesis, KTH Royal Institute of Technology.
- Moore, K.L. (1993). *Iterative Learning Control for Deterministic Systems*. Springer-Verlag.
- Oomen, T. (2020). Learning for Advanced Motion Control. In *Int. Work. Adv. Motion Control*, 65–72. Institute of Electrical and Electronics Engineers, United States.
- Poot, M., Portegies, J., and Oomen, T. (2021). Kernel-Based Learning Control for Iteration-Varying Tasks Applied to a Printer With Friction. In *IEEE/ASME Int. Conf. Adv. Intell. Mechatronics*, 1052–1057.
- Schoukens, J. and Ljung, L. (2019). Nonlinear System Identification: A User-Oriented Road Map. *IEEE Control Syst.*, 39(6), 28–99.
- Söderström, T. and Stoica, P. (1989). *System Identification*. Prentice-Hall Software Series. Prentice Hall.
- van de Wijdeven, J. and Bosgra, O.H. (2010). Using basis functions in iterative learning control: analysis and design theory. *Int. J. Control*, 83 (4), 661–675.
- van Haren, M., Poot, M., Portegies, J., and Oomen, T. (2022). Position-Dependent Snap Feedforward: A Gaussian Process Framework. *arXiv*.
- van Herpen, R., Oomen, T., and Steinbuch, M. (2014). Optimally conditioned instrumental variable approach for frequency-domain system identification. *Automatica*, 50(9), 2281–2293.
- van Zundert, J., Bolder, J., and Oomen, T. (2016). Optimality and flexibility in Iterative Learning Control for varying tasks. *Automatica*, 67, 295–302.
- van Zundert, J. and Oomen, T. (2018). On inversion-based approaches for feedforward and ILC. *Mechatronics*, 50, 282–291.