Iterative Adaptive Control: Windsurfing with Confidence *

Robert L. Kosut[1,2]

[1] SC Solutions, Inc., 3211 Scott Blvd., Santa Clara, CA 95054, USA
kosut@scsolutions.com
[2] Research supported by DARPA, Applied Computation & Mathematics Program, under NASA Grant No. NAG-1-1964.

1 Introduction

SEEKER: How do I gain Good Judgement?
TEACHER: You must first acquire Wisdom.
SEEKER: How do I gain Wisdom?
TEACHER: From Bad Judgement.

Windsurfing, Learning, and Unfalsification

Learning, apparently, is the result of making mistakes. A good example is learning to windsurf. The approach I took is simple. Fall Often. This is the approach I seem to use regularly when learning anything. It helped a great deal that I did this in Tahiti where both the air and water were 80 degrees and there was little wind. I was in Tahiti following my first visit with Brian Anderson in Canberra in 1984. A few years later when I was describing my learning experience at windsurfing to Brian, who by the way is an excellent windsurfer, we realized that how a human learns windsurfing might be a procedure that could be used for adaptive control. We speculated about human learning of physical activities as follows. The human first learns to control over a limited bandwidth and learning pushes out the bandwidth over which an accurate knowledge of the "plant" is known. The human first implements a low-gain controller and learning allows the loop to be tightened. The human tends to "over-react" when the environment changes, often resorting to high-gain or no-gain. Then falls. The few instants just before falling may provide some new information about the next higher increments in the bandwidth of the system. So of course mistakes are fundamental to learning. (This is why stabilizing feedback is *negative*; the popular concept of *positive* feedback

* A chapter in *Model Identification and Adaptive Control: From Windsurfing to Telecommunications*, G. Goodwin (ed.), Springer-Verlag 2001, a collection of contributions written in honour of Professor Brian Anderson's 60th birthday.

may be useful, in fact, when it is actually negative!) Our speculations on this "windsurfing" approach to adaptive control was first expressed in [2] and subsequently filled out in a series of papers [24,25].

In general, the learning process has the following steps: make observations, construct a model, and validate the model against new data. The last step means the model is used for some purpose and tested against reality in some way. If the intended purpose of the model is control design, then the penultimate validation of the model is to *implement* a controller based on the model. Of course the implementation step generates new data, and this is usually where new mistakes are uncovered. For example, just prior to "falling."

Validation is perhaps a misnomer. One can never prove that a model will be able to accurately predict the future. More precisely, the data can *falsify* a model, *i.e.*, the model may prove to be incapable of fully explaining the data. This step, that is, discarding bad models, could be done before implementation. To emphasize this kind of in-validation, we use the more precise, but awkward term: *unfalsification*. The essence of unfalsification is this: we use data from one experiment and ask if it could have come from another experiment. This second experiment need not be performed because if it is shown that the data is not consistent with the second "thought experiment," then there is no need to perform it. If it cannot be proven inconsistent, then the only recourse is to do the second experiment, *e.g.*, implement the controller under consideration. No guarantees are offered other than it fails to be proven inconsistent.

In order to relate this to control design it is natural to think in terms of a model based design, *i.e.*, design the control using a model of the system being controlled. However, no mathematical model of a physical system is ever exact. Hence, a complete model for purposes of control design should include both the nominal model as well as an explicit description of the uncertainty between the model and the physical system. Such a model is referred to as an *uncertainty model* and is the starting point for robust control design [40]. Typically this uncertainty characterization is a guess, although an informed one. Inferring the uncertainty from data alone is not provided by a straight forward application of the classical method of indirect adaptive control *i.e.*, system identification followed by control design. Classical system identification alone does not offer this feature, *e.g.*, [27]. Recent efforts, however, which modify the classical approach to system identification, specifically for control design, do address this problem, *e.g.*, [28].

Classical Adaptive Control

Classical adaptive control methods, typified by the "indirect" scheme shown in figure 1, can work well in many instances, *e.g.*, when the plant variations are slow and the input excitation is concentrated at those frequencies which are critical. But even in this almost ideal case, initial plant-model mismatch

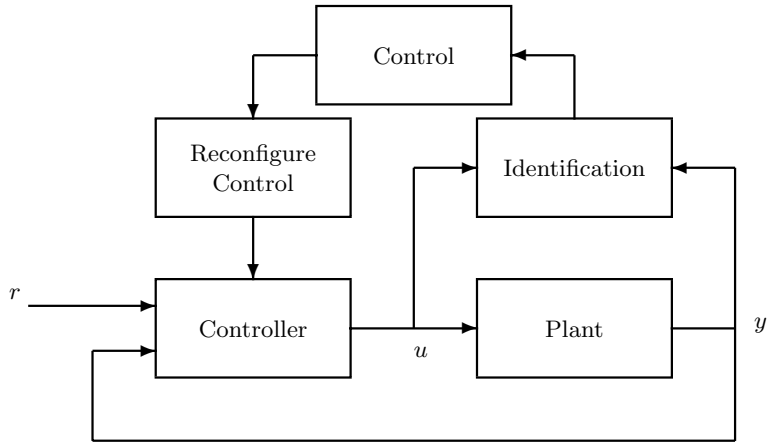

Fig. 1. Classical indirect adaptive control: the control parameters are adjusted "indirectly" via identification and control re-design based on the (model) estimate.

can induce unacceptably large transients. In more volatile circumstances the adaptive system can exhibit instabilities, limit cycles, and chaotic behavior, even when initialized at a desired controller parameter setting. Theoretical studies have exposed some of the underlying mechanisms for stability and instability of these classical adaptive control approaches, *e.g.*, [3], [1]. A rough explanation is that these classical adaptive schemes require gradients of the error, which depend of course on the unknown plant, and hence, have to be approximated. (There are some recently developed methods which provide "unbiased" estimates of the gradients from data, *e.g.*, [12], but they require specialized experiments and thus are not al all suited to handling rapid or unexpected system variations.) The same problems plague the "direct" adaptive schemes where controller parameters are directly adjusted to minimize the error between the adaptively controlled plant and a reference system.

Unfasified Adaptive Control

Over the past several years a new approach to adaptive control, based on the paradigm of "unfalsification," has been under investigation. The approach is based on a new perspective to the adaptive control problem and is important in two ways: (1) as an "indirect" adaptive robust control method it provides a model error bound from the experimental data which can be used to assess the model quality for robust and/or fault tolerant control design, and (2) as a "direct" adaptive control method it can be used to eliminate, *without further implementation*, those controllers that would not satisfy the performance goals. These approaches appear to provide the theoretical foundation

for synthesizing reliable adaptive control systems which are not prone to the difficulties encountered by classical approaches.

Unfalsification provide a means to estimate an uncertainty model from finite data which includes both disturbance and dynamic uncertainty as well as the usual parameters in the standard prediction error models. In this paper we will show how unfalsification can be used for *direct iterative adaptive control*, *i.e.*, the control parameters are adjusted directly. In our earlier work on iterative (windsurfing) adaptive control [24,25], tests for unfalsification were not available.

Uncertainty model unfalsification using finite time domain data, which is the underlying basis for the work presented in this paper, was first described by Poolla *et al.*[33]. Earlier work using frequency domain data was presented in [35], and some precursors to unfalsification can be found in [16–18]. Further extensions and applications to plant uncertainty model unfalsification can be found in [19], [20], [21], and [26]. A method employing a probabilistic description of dynamic uncertainty is given in [10]. The origin of the ideas for direct controller unfalsification are presented by Safonov and Tsao in [34] and the references therein. The mathematical basis for unfalsification of linear-time-invariant systems can be found in [11] and [8]. Computations using convex programming is discussed in [37,38].

2 Iterative Adaptive Control

The roots of iterative adaptive control can be traced to the dual control concept (see, *e.g.*, [3, Ch. 7]), which typically involves *indirect adaptation*, *i.e.*, identification followed by control parameter adjustment. A survey of iterative identification and control schemes is given in [9]. Of particular relevance to the work presented here – for purposes of comparison – is [39], [5], and [4] which describe how data filters can be selected to make the identification and control criteria merge; the windsurfer approach to adaptation and learning, as described in [24,25], where the closed-loop bandwidth is gradually increased every iteration; and [12], which describes a direct iterative controller design method.

A generic iterative adaptive control system is depicted in figure 2. The adaptive part of the controller consists of a parameter estimator and a control design algorithm connected in series through a sample and hold. The latter is what makes the system "iterative." That is, the next controller design is based on data collected while the previous controller was in place.

The system consists of two feedback "loops" each operating at different sampling rates. The inner loop, operating at the fast rate, consists of the plant and controller, where u is the control input to the plant, y is the sensed output from the plant, and r is the reference command to the controller. The outer loop, operating at the slow rate, consists of the plant parameter estimator and control parameter design, The sequence of parameter estimates, $\bar{\theta}$, are

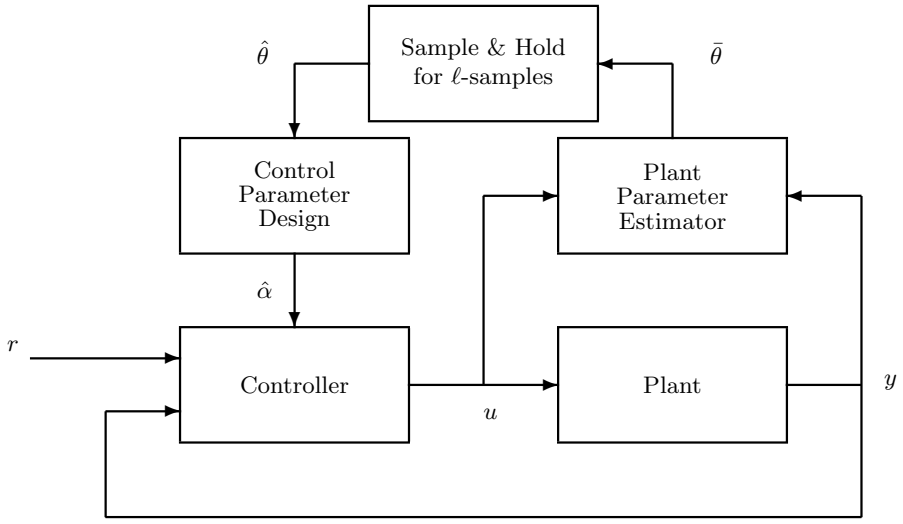

Fig. 2. Iterative adaptive control system.

produced at the end of every data collection interval of ℓ-samples, and hence, depend on the prior applied sequence of controller parameters, $\hat{\alpha}$, which are based on $\hat{\theta}$, the prior plant parameter sequence, and so on. Thus, $\hat{\theta}$ and $\hat{\alpha}$, are piece-wise constant vector sequences, *i.e.*, constant over every ℓ-samples. Specifically, during the i-th iteration (ℓ-data collection interval), that is, for $t = 1 + (i-1)\ell, \ldots, i\ell$, let $\theta_\ell^i \in \mathbf{R}^p$ denote the plant parameters and let $\alpha_\ell^i \in \mathbf{R}^q$ denote the corresponding control parameters. The relation between θ_ℓ^i and α_ℓ^i is typically algebraic and depends on the design procedure, *i.e.*,

$$\alpha_\ell^i = \kappa(\theta_\ell^i) \tag{1}$$

Let $(y^i, u^i) \in \mathbf{S}^\ell \times \mathbf{S}^\ell$ denote[1] the plant output and input data recorded during the i-th data collection period, *i.e.*,

$$\begin{aligned} y^i &= \{\ y_{1+(i-1)}, \ldots, y_{i\ell}\ \} \\ u^i &= \{\ u_{1+(i-1)}, \ldots, u_{i\ell}\ \} \end{aligned} \tag{2}$$

[1] $\mathbf{S}^\ell$ denotes the set of scalar sequences of length ℓ, *i.e.*, $x \in \mathbf{S}^\ell \leftrightarrow x = \{\ x_t \mid t \in [1, \ell]\ \}$ where t are the uniformly spaced integer samples.

After every data collection period of ℓ samples, a new parameter estimate, denoted by θ_ℓ^{i+1}, is determined by solving an optimization problem of the form[2],

$$\theta_\ell^{i+1} = \arg \min_{\theta \in \mathbf{\Theta}} V_\ell^i(\theta) \tag{3}$$

where $\mathbf{\Theta} \subset \mathbf{R}^p$ is the set of possible parameters, which in most cases is simply $\mathbf{R}^p$. The objective function, $V_\ell^i(\theta)$, depends on the data $(y^i, u^i) \in \mathbf{S}^\ell \times \mathbf{S}^\ell$, as denoted by the superscript i and the subscript ℓ, and on θ, where the precise dependence is determined by the estimation and/or control design criterion. For "indirect adaptive" schemes, the criterion is related to providing a good fit to the data based on an assumed uncertainty model for the plant system. Hence, $\alpha = \kappa(\theta)$ simply denotes how the plant uncertainty model parameters relate to the corresponding control design parameters. For the "direct adaptive" schemes, as the name implies, the control parameter is itself adjusted, *i.e.*, $\alpha \equiv \theta$, and the criterion is directly related to closed-loop performance. Actually even in the indirect case, the criterion is constructed to be as close as possible to a closed-loop performance criterion or at least useful for control design.

2.1 Convergence

Convergence analysis is very difficult because the data $(y^i, u^i) \in \mathbf{S}^\ell \times \mathbf{S}^\ell$ depends on *all* past control parameter switching $\{\ \alpha_\ell^1, \ldots, \alpha_\ell^i\ \}$, which in turn depend on all past plant parameter estimates $\{\ \theta_\ell^1, \ldots, \theta_\ell^i\ \}$. However, for long data collection periods, provided that all past controllers are stabilizing, the system memory of past controllers fades. Hence, in the limit, with infinite data collected during every iteration, the data collected during the i-th interval only depends on the last parameter values (α^i, θ^i), *i.e.*,

$$\lim_{\ell \to \infty} y^i = y(\alpha^i), \quad \lim_{\ell \to \infty} u^i = u(\alpha^i), \quad \alpha^i = \kappa(\theta^i) \tag{4}$$

Similarly, the objective function becomes,

$$\lim_{\ell \to \infty} V_\ell^i(\theta) = V(\theta, \alpha^i), \quad \alpha^i = \kappa(\theta^i) \tag{5}$$

So in the limiting case of infinite data, the parameter estimation step at iteration i is,

$$\theta^{i+1} = \arg \min_{\theta \in \mathbf{\Theta}} V(\theta, \alpha^i), \quad \alpha^i = \kappa(\theta^i) \tag{6}$$

[2] The optimization operation "arg min" as used here and throughout the paper is to be understood to mean the minimizing argument, or if there is no unique minimum, then "arg min" refers to the set of (local) minima, *i.e.*, $\arg\min f(x) = \{\ x \mid f_x(x) = 0,\ f_{xx}(x) > 0\ \}$.

As observed in [13], for infinite data, convergent parameter values are equivalently fixed-points of the mapping $\Gamma : \mathbf{R}^p \mapsto \mathbf{R}^p$, where

$$\Gamma(\theta) = \arg \min_{\psi \in \mathbf{\Theta}} V(\psi, \alpha) \;\text{ subject to }\; \alpha = \kappa(\theta) \tag{7}$$

Hence, if $\hat{\theta}$ is a fixed-point, then

$$\hat{\theta} = \arg \min_{\psi \in \mathbf{\Theta}} V(\psi, \hat{\alpha}), \quad \hat{\alpha} = \kappa(\hat{\theta}) \tag{8}$$

and must satisfy the necessary condition for optimality, namely,

$$\frac{\partial}{\partial \theta} V(\theta, \kappa(\hat{\theta}))\bigg|_{\theta = \hat{\theta}} = 0 \tag{9}$$

However, the minimizer, θ_{opt}, of $V(\theta, \kappa(\theta))$ satisfies,

$$\frac{\partial}{\partial \theta} V(\theta, \kappa(\theta))\bigg|_{\theta = \theta_{opt}} = 0 \tag{10}$$

As pointed out in [13], the fixed-point $\hat{\theta}$ is not likely to be the same as θ_{opt}. So it appears that iterative schemes have a built-in flaw. Even if the infinite-data estimation criterion, $V(\theta, \kappa(\theta))$, is constructed to be a sensible control criterion, no iterative algorithm can be guaranteed to reach the minimum, or at least a local minimum. This has led some researchers to seek another path, *e.g.*, in [12] the authors show how to obtain an unbiased estimate of the gradient (and Hessian) of error signals with respect to control parameters by performing a series of specialized experiments.

2.2 Relation to Slow Adaptation

This "flawed" property of iterative adaptation is a recrudescence of the identical property of all slowly varying parameter adaptive algorithms, *e.g.*, see [1], [32], or the chapter on averaging analysis in [3]. As an illustration, consider the adaptive algorithm,

$$\begin{aligned} \theta_{t+1} &= \theta_t + \mu \phi_t(\alpha_t) \varepsilon_t(\theta_t, \alpha_t) \\ \varepsilon_t(\theta, \alpha) &= y_t(\alpha) - \phi_t(\alpha)^T \theta_t \\ \alpha &= \kappa(\theta) \end{aligned} \tag{11}$$

where $\mu > 0$ adjusts the speed of adaptation and where $\varepsilon_t(\theta_t, \alpha_t)$ is the prediction error associated with an ARX plant model being identified in closed-loop with the control parameters set to α_t. If the adaptation is slow, *i.e.*, μ is sufficiently small, then as shown in [1], there exists $(\theta_t^\mu, \alpha_t^\mu)$ satisfying (11) in an order-μ neighborhood of $(\hat{\theta}, \hat{\alpha})$ obtained from[3] ,

$$\Big\langle \phi(\hat{\alpha}) \varepsilon(\hat{\theta}, \hat{\alpha}) \Big\rangle = 0, \quad \hat{\alpha} = \kappa(\hat{\theta}) \tag{12}$$

[3] $\Big\langle x \Big\rangle$ denotes time average, *i.e.*, $\Big\langle x \Big\rangle = \lim_{\ell \to \infty} \frac{1}{\ell} \sum_{t=1}^{\ell} x_t$

Equivalently,

$$\begin{aligned}\hat{\theta} &= \Big\langle \phi(\hat{\alpha})\phi(\hat{\alpha})^T \Big\rangle^{-1} \Big\langle \phi(\hat{\alpha}) y(\hat{\alpha}) \Big\rangle \\ \hat{\alpha} &= \kappa(\hat{\theta})\end{aligned} \tag{13}$$

This is precisely the solution to (9) if $V(\theta, \alpha)$ is the standard RMS (least-squares) identification criterion,

$$V(\theta, \alpha) = \Big\langle \varepsilon(\theta, \alpha)^2 \Big\rangle \tag{14}$$

In this case the reason that $(\hat{\theta}, \hat{\alpha})$ is not the optimum solution given by (10) is that

$$\phi(\kappa(\theta)) \neq \frac{\partial}{\partial \theta} \varepsilon(\theta, \kappa(\theta)) \tag{15}$$

The conditions which insure that $(\theta_t^\mu, \alpha_t^\mu)$ satisfies (11) for small μ do not insure that the solution is stable. Further conditions are required [1], [32]. Similar results hold for "direct" adaptive algorithms, *e.g.*,

$$\begin{aligned}\alpha_{t+1} &= \alpha_t + \mu \psi_t(\alpha_t) \varepsilon_t(\alpha_t) \\ \varepsilon_t(\alpha) &= y_t(\alpha) - y_t^{\text{ref}} \\ u_t &= \psi_t(\alpha_t)^T \alpha_t\end{aligned} \tag{16}$$

In this case for sufficiently small μ, there is a solution α_t^μ in an order-μ neighborhood of $\hat{\alpha}$ satisfying

$$\Big\langle \psi(\hat{\alpha}) \varepsilon(\hat{\alpha}) \Big\rangle = 0 \tag{17}$$

As in the previous case, $\hat{\alpha}$ is not the solution to (9) with the logical choice,

$$V(\alpha) = \Big\langle \varepsilon(\alpha)^2 \Big\rangle \tag{18}$$

because

$$\psi(\alpha) \neq \frac{\partial}{\partial \alpha} \varepsilon(\alpha) \tag{19}$$

But all is not lost for iterative schemes. By introducing an unfalsification step we can develop iterative schemes which avoid the above "convergence" problems. Perhaps convergence is not an appropriate concept when conclusions are being drawn, in this case eliminated, strictly from finite data. Let's leave this to the philosophers.

3 Parametrization and Performance

3.1 Parametrization

We will assume throughout that the control is given by,

$$u = C(\alpha)(r - y) \tag{20}$$

where for fixed $\alpha \in \mathbf{R}^q$, $C(\alpha^i) \in \mathbf{LTI}$[4] The controller parameters, α, are to be adaptively adjusted. There are of course many possible ways to parametrize the controller. For example, the control parameters, α, can consist of all the numerator and denominator transfer function coefficients up to a specified degree, thereby restricting the controller order. For example, all order-n controllers can be parametrized as follows:[5]

$$\begin{aligned} C(\alpha) &= N(\alpha)D(\alpha)^{-1} \\ N(\alpha) &= b_0 + b_1 z^{-1} + \cdots + b_n z^{-n} \\ D(\alpha) &= 1 + a_1 z^{-1} + \cdots + a_n z^{-n} \\ \alpha &= [a_1 \cdots a_n \; b_0 \cdots b_n]^T \in \mathbf{R}^{2n+1} \end{aligned} \tag{21}$$

Another paramtrization is all the PI controllers, *i.e.*,

$$C(\alpha) = \alpha_P + \alpha_I \frac{z^{-1}}{1 - z^{-1}}, \quad \alpha = [\alpha_P \; \alpha_I \;]^T \in \mathbf{R}^2 \tag{22}$$

Obviously many parametrizations are possible, *e.g.*, co-prime factorizations where the factors are affine in the control parameters, or direct affine parametrizations with orthogonal basis functions [36].

There are a number of other interesting possibilities. For example, a list of controllers each corresponding to some measurable condition of the system. This approach is somewhat like "multiple model control" or "gain-scheduling." Control structures could have free parameters (α) corresponding to a prescribed set of anticipated faults, *e.g.*, changes in sensor bias and scale factors, changes in actuator gains and bandwidth, and changes in a mechanical structure resulting in spillover of elastic modes into the closed-loop bandwidth. A similar approach was investigated in [7] for a switching missile autopilot.

3.2 Performance

We would like the closed-loop system to behave like the reference system,

$$y_{\text{ref}} = T_{\text{ref}} \; r \tag{23}$$

for a specified system $T_{\text{ref}} \in \mathbf{LTI}$. "Behave like" can have a variety of meanings. For example, it could mean that the output error, $y - T_{\text{ref}} \; r$, should be

[4] **LTI** is the set of linear-time-invariant systems with rational transfer functions.
[5] z^{-k} denotes the k-delay operator.

small relative to the size of the command r. An example of such a specification is that[6] ,

$$\big\|y - T_{\rm ref}\ r\big\|_{rms} \le \rho \big\|r\big\|_{rms}, \quad \forall \big\|r\big\|_{rms} < \infty \tag{24}$$

In this case, we are not looking for a response to a *specific* r, such as a sinusoid at one frequency or a step, rather, for every possible r such that $\big\|r\big\|_{rms} < \infty$. For example, if the plant system is given by,

$$y = Pu \tag{25}$$

where $P \in \mathbf{LTI}$, and the controller is given by,

$$u = C(r - y) \tag{26}$$

with $C \in \mathbf{LTI}$, then (24) is equivalent to,

$$\|T(P,C) - T_{\rm ref}\|_{\mathbf{H}_\infty} \le \rho \tag{27}$$

with $T(P,C)$ given by,

$$T(P,C) = (1 + PC)^{-1}PC \tag{28}$$

This is clearly a measure of the error between the transfer functions of the closed-loop and reference systems. Hence, for a given controller parametrization, $C(\alpha) \in \mathbf{LTI}$, and a given plant parametric model, $P(\theta) \in \mathbf{LTI}$, the optimal control design (1) is,

$$\kappa(\theta) = \arg\min_{\alpha} \|T(P(\theta), C(\alpha)) - T_{\rm ref}\|_{\mathbf{H}_\infty} \tag{29}$$

Hence, the *designed* closed-loop system $T(P(\theta), C(\kappa(\theta))$ is the closest $\mathbf{H}_\infty$ approximation to $T_{\rm ref}$, the reference system. Since it is unrealistic to expect that the plant is in the model set, *i.e.*, $P \neq P(\theta)$, it follows that the actual, or *achieved* closed-loop system, $T(P, C(\kappa(\theta))$, may be quite different. Of course the most desirable goal is,

$$\alpha_{opt} = \arg\min_{\alpha} \|T(P, C(\alpha)) - T_{\rm ref}\|_{\mathbf{H}_\infty} \tag{30}$$

The performance measure can be modified to penalize control activity, *e.g.*,

$$\big\|y - T_{\rm ref}\ r\big\|^2_{rms} + \lambda \big\|u\big\|^2_{rms} \le \rho^2 \big\|r\big\|^2_{rms}, \quad \forall \big\|r\big\|_{rms} < \infty \tag{31}$$

Again, if the plant and controller are LTI, then this measure is equivalent to,

$$\left\| \begin{bmatrix} T(P,C) - T_{\rm ref} \\ \lambda Q(P,C) \end{bmatrix} \right\|_{\mathbf{H}_\infty} \le \rho \tag{32}$$

where

$$Q(P,C) = (1 + PC)^{-1}C \tag{33}$$

[6] The RMS-norm of a sequence (technically a semi-norm) is defined as $\big\|x\big\|_{rms} = \left(\lim_{\ell \to \infty} \frac{1}{\ell} \sum_{t=1}^{\ell} x_t^2 \right)^{1/2}$

4 Uncertainty Model Unfalsification

4.1 Unfalsification

The generic uncertainty model unfalsification problem is as follows:

Given scalar data sequences $e, v \in \mathbf{S}^\ell$, establish necessary and sufficient conditions for the existence of a disturbance sequence $w \in \mathbf{S}^\ell$ and a causal system Δ such that

$$w \in \mathbf{W}(\sigma), \quad \Delta \in \mathbf{\Delta}(\delta) \tag{34}$$

and which are consistent with the model

$$e_t = w_t + (\Delta v)_t, \quad t \in [1, \ell] \tag{35}$$

The sets $\mathbf{W}(\sigma)$ and $\mathbf{\Delta}(\delta)$ denote, respectively, a set of sequences with norm bounded by σ and a set of systems with gain bounded by δ.

4.2 Uncertainty Model Forms

The data sequence e is often obtained as the *prediction error* associated with an assumed model of the system, and v is a function of other sensed signals, the choice reflecting the type of dynamic uncertainty, or model error. For example, consider the standard prediction error form in [27],

$$e = H^{-1}(y - Pu), \qquad P, H \in \mathbf{LTI} \tag{36}$$

If $v = u$, then Δ represents *additive* model error, *i.e.*, the uncertainty model set is given by,

$$\mathbf{J}(\sigma, \delta) = \{\ y, u\ \mid y = Pu + H\Delta u + Hw, \quad w \in \mathbf{W}(\sigma), \quad \Delta \in \mathbf{\Delta}(\delta)\ \}$$

If $v = Gu$, then Δ represents *multiplicative* model error, *i.e.*, the uncertainty model set becomes,

$$\mathbf{J}(\sigma, \delta) = \{\ y, u\ \mid y = Pu + H\Delta Gu + Hw, \quad w \in \mathbf{W}(\sigma), \quad \Delta \in \mathbf{\Delta}(\delta)\ \}$$

There are clearly many variations one could include, *e.g.*, combinations of additive and multiplicative model errors, co-prime factor uncertainty, and so on, ultimately leading to the uncertainty structures described by the more inclusive linear fractional representation familiar in robust control design, *e.g.*, [30]. In addition, the error could be obtained from a *parametric* prediction error model with parameters associated with transfer function coefficients which characterize the input/output and disturbance dynamics, *i.e.*,

$$e(\theta) = H(\theta)^{-1}\,(y - P(\theta)u) \tag{37}$$

with $\theta \in \mathbf{\Theta}$, the set of parameters for which the predictor is stable [27].

4.3 Disturbance Uncertainty

There are many ways to characterize the disturbance set $\mathbf{W}(\sigma)$. For example, consider the following sets of finite sequences:

- **Rms-bounded noise**

$$\mathbf{W}_{\rm rms}(\sigma) = \left\{ \; w \in \mathbf{S}^\ell \; \middle| \; \frac{1}{\ell}\big\|w\big\|^2 \le \sigma^2 \; \right\} \tag{38}$$

- **Time-domain white noise** [31]

$$\mathbf{W}_{\rm wht\_time}(\gamma, m) = \left\{ \; w \in \mathbf{S}^\ell \; \middle| \; |r_w(\tau)| \le \gamma r_w(0) \; \right\} \tag{39}$$

  where $r_w(\tau)$ is the auto-correlation of w,

$$r_w(\tau) = \frac{1}{\ell}\sum_{t=1}^{\ell-\tau} w_t w_{t+\tau}, \quad \tau \in [0, m-1] \le \ell \tag{40}$$

  Observe that $r_w(0) = \big\|w\big\|^2/\ell$.
- **Frequency-domain white noise** [29]

$$\mathbf{W}_{\rm wht\_freq}(\sigma, \epsilon, m) = \left\{ \; w \in \mathbf{S}^\ell \; \middle| \; |eig\left\{R_m(w)\right\}/\sigma^2 - 1| \le \epsilon \; \right\} \tag{41}$$

  where

$$R_m(w) = \begin{bmatrix} r_w(0) & \cdots & r_w(m-1) \\ \vdots & \ddots & \vdots \\ r_w(m-1) & \cdots & r_w(0) \end{bmatrix} \tag{42}$$

The disturbance set $\mathbf{W}_{\rm rms}(\sigma)$ is the simplest of choices for deterministically characterizing "noise." The main advantage is that it is a convex set and therefor easy to handle in optimization. However, there are no restrictions preventing correlation with inputs and so the "worst-case" can occur. As shown above, characterizations of deterministic sets which resemble white noise have been examined in [29] in the frequency domain with application to system identification and in [31] for both time and frequency domains with application to robust control. The set $\mathbf{W}_{\rm wht\_time}(\gamma, m)$ is essentially one of the standard white noise test where γ is chosen from χ^2 distribution tables; m is the *lag window* used to smooth the correlation function. The set $\mathbf{W}_{\rm wht\_freq}(\sigma, \epsilon, m)$ is shown in [29] to also be useful for white noise testing; m again is the lag window, σ^2 is the rms-level of w and hence, the average level of the spectrum of w, and $\epsilon \in (0, 1)$ determines the "flatness" of the spectrum. Clearly these latter sets do preserve the character of white noise, but they are not convex. However, they are no worse than quadratic and so may be quite amenable to conjugate-gradient methods of optimization. The work reported in [23] shows a two-step procedure involving a Kalman filter for unfalsifying stochastic disturbance signals.

4.4 Gain-Bounded Dynamic Uncertainty

Uncertain dynamics can also be characterized in a number of ways. Consider the following gain-bounded, time-invariant (**TI**) dynamic uncertainty sets:

- **Linear (LTI)**

$$\mathbf{\Delta}_{\mathbf{L}TI}(\delta) = \{\ \Delta \in \mathbf{LTI}\ \mid\ \|\Delta v\|_{rms} \le \delta \|v\|_{rms},\ \ \forall \|v\|_{rms} < \infty\ \} \quad (43)$$

  Since $\Delta \in \mathbf{LTI}$, the gain bound condition is equivalent to the frequency domain bound:

$$|\Delta(e^{j\omega})| \le \delta,\ \ \omega \in [-\pi, \pi] \quad (44)$$

- **Incrementally nonlinear (INTI)**

$$\mathbf{\Delta}_{\mathrm{INTI}} = \{\ \Delta \in \mathbf{TI}\ \mid\ \|\Delta v_1 - \Delta v_2\|_{rms} \le \delta \|v_1 - v_2\|_{rms},\ \ \forall \|v_1\|_{rms}, \|v_2\|_{rms} < \infty\ \} \quad (45)$$

- **Nonlinear (NTI)**

$$\mathbf{\Delta}_{\mathrm{NTI}} = \{\ \Delta \in \mathbf{TI}\ \mid\ \|\Delta v\|_{rms} \le \delta \|v\|_{rms},\ \ \forall \|v\|_{rms} < \infty\ \} \quad (46)$$

4.5 Unfalsification

Consequences of unfalsification are summarized in the following.

(i) **Finite-Data Test**
Given data sequences $e, v \in \mathbf{S}^\ell$, there exists a sequence $w \in \mathbf{S}^\ell$ and a causal system Δ such that,

$$e_t = w_t + (\Delta v)_t,\ \ t \in [1, \ell] \quad (47)$$

with $w \in \mathbf{W}_{\mathrm{rms}}(\sigma)$ if and only if

$$\frac{1}{\ell}\|w\|^2 \le \sigma^2 \quad (48)$$

and such that:

- *$\Delta \in \mathbf{\Delta}_{\mathbf{LTI}}(\delta)$ if and only if,*

$$\mathcal{T}\{e-w\}^T \mathcal{T}\{e-w\} - \delta^2 \mathcal{T}\{v\}^T \mathcal{T}\{v\} \le 0 \quad (49)$$

  with $(\mathcal{T}\{e\}, \mathcal{T}\{v\}, \mathcal{T}\{w\})$ the $\ell \times \ell$ Toeplitz matrices formed from the sequences (e, v, w), respectively, e.g.,

$$\mathcal{T}\{e\} = \begin{bmatrix} e_1 & 0 & \cdots & 0 \\ e_2 & e_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ e_\ell & e_{\ell-1} & \cdots & e_1 \end{bmatrix}$$

- *$\Delta \in \mathbf{\Delta}_{\mathbf{INTI}}(\delta)$ if and only if, $\forall m - n = 0 : \ell$ and $\forall t \in [1, \ell]$,*

$$\|(z^n - z^m)(e - w)\|_{\mathbf{L}_2[1,t]} \le \delta \|(z^n - z^m)v\|_{\mathbf{L}_2[1,t]} \quad (50)$$

  where z^k is the k-forward shift operator, i.e., if $x = \{\ x_1, x_2, \dots\ \}$ then $z^k x = \{\ 0, \dots, 0, x_1, x_2, \dots\ \}$ with k-zeros.

- $\Delta \in \mathbf{\Delta_{NTI}}(\delta)$ *if and only if* $\forall t \in [1,\ell]$,
$$\big\|e-w\big\|_{\mathbf{L}_2[1,t]} \le \delta \big\|v\big\|_{\mathbf{L}_2[1,t]} \tag{51}$$

(ii) **Uncertainty Tradeoff**
The (σ,δ) *boundary between falsified and unfalsified uncertainty models for a given finite data set with* $w \in \mathbf{W}_{\rm rms}$ *and* $\Delta \in \mathbf{\Delta}_\mu(\delta)$, $\mu = \mathbf{LTI},\ \mathbf{INTI},\ \mathbf{NTI}$ *is determined by solving:*

$$\sigma_\mu(\delta) = \min\Big\{\ \ell^{-1}\big\|e-\Delta v\big\|_{\mathbf{L}_2[1,\ell]}\ \mid \Delta \in \mathbf{\Delta}_\mu(\delta)\ \Big\}, \quad 0 \le \delta \le \delta_\mu \tag{52}$$

with

$$\delta_\mu = \min\{\ \delta\ \mid e_t = (\Delta v)_t,\ t\in[1,\ell],\ \Delta \in \mathbf{\Delta}_\mu(\delta)\ \} \tag{53}$$

(iii) **Nesting**
For all $\delta \ge 0$,

$$\sigma_{\mathbf{NTI}}(\delta) \le \sigma_{\mathbf{INTI}}(\delta) \le \sigma_{\mathbf{LTI}}(\delta) \tag{54}$$

and

$$\delta_{\mathbf{NTI}} < \delta_{\mathbf{INTI}} < \delta_{\mathbf{LTI}} \tag{55}$$

Comments

(1) The results in part (i) for $\Delta \in \mathbf{\Delta_{LTI}}(\delta)$ and the necessity for $\Delta \in \mathbf{\Delta_{NTI}}(\delta)$ (which is the same as the necessary and sufficient conditions for gain-bounded linear-time-varying (LTV) systems) is found in [33]. Proof of the remaining results in (i) can be found in [21].

(2) The tradeoff and nesting results follow from convexity of the uncertainty sets. All the results can be extended when the error is formed from the ARX parametric prediction error model with efficient computations using LMIs, [6]. Output error and other linear fractional parameter forms are not convex sets, and the nesting and tradeoff results are thus not guaranteed. A further discussion of the relation of NTI and LTI uncertainty sets, can be found in [19–21].

5 Controller Unfalsification

In a series of papers by Safonov *et al.*(see [34] and the references therein), it is shown how to *directly* falsify a candidate controller before it is implemented. The procedure for controller unfalsification is essentially the same as that for uncertainty model unfalsification, but applied to the closed-loop specification. Specifically, the closed-loop specification (24) can be viewed as the uncertainty model set,

$$\mathbf{J}(\rho) = \Big\{\ y,r\ \Big|\ \big\|y - T_{\rm ref}\ r\big\|_{rms} \le \rho\big\|r\big\|_{\rm rms}\ ,\ \ \forall\big\|r\big\|_{rms} < \infty\ \Big\} \tag{56}$$

The goal is to adjust the parameters $\alpha \in \mathbf{R}^q$ such that the controller,

$$u = C(\alpha)(r-y) \tag{57}$$

makes ρ as small as possible. From the previous discussion on unfalsification of uncertainty models, without any further assumptions about the plant system, the specification set is equivalently expressed as,

$$\mathbf{J}(\rho) = \{\ y, r\ \mid y - T_{\text{ref}}\ r = \Delta r, \quad \Delta \in \mathbf{\Delta_{NTI}}(\rho)\ \} \tag{58}$$

If we make the additional assumption that the closed-loop system is LTI, then the specification becomes the LTI uncertainty set,

$$\mathbf{J}(\rho) = \left\{\ y, r\ \middle|\ y - T_{\text{ref}}\ r = \Delta r, \quad \|\Delta\|_{\mathbf{H}_\infty} \leq \rho\ \right\} \tag{59}$$

The finite data tests provided in the previous section can now be used to compute the performance ρ. Let $(y, u, r) \in \mathbf{S}^\ell \times \mathbf{S}^\ell \times \mathbf{S}^\ell$ denote the measured data for $t \in [1, \ell]$, where ℓ is essentially the current time. If $\Delta \in \mathbf{NTI}$, then it follows from (51) that the *measured* performance is,

$$\begin{aligned} \rho^\ell_{meas} &= \min \left\{\ \rho\ \middle|\ \|\varepsilon\|_{\mathbf{L}_2[1,t]} \leq \rho \|r\|_{\mathbf{L}_2[1,t]}\ , \quad t \in [1, \ell]\ \right\} \\ \varepsilon &= y - T_{\text{ref}}\ r \end{aligned} \tag{60}$$

If $\Delta \in \mathbf{LTI}$, then it follows from (49) that the *measured* performance is,

$$\rho^\ell_{\text{meas}} = \min \left\{\ \rho\ \middle|\ \mathcal{T}\{\varepsilon\}^T\, \mathcal{T}\{\varepsilon\} - \rho^2 \mathcal{T}\{r\}^T\, \mathcal{T}\{r\} \leq 0\ \right\} \tag{61}$$

Both of these measure the performance of the already implemented controller, which may have already switched (adapted) several times. The question arises: could the existing data record be informative about an untried *candidate controller*, say $C(\alpha)$, whose predicted performance might be better than the measured performance? To answer the question, consider, the following "thought experiment:"

If a candidate controller $C(\alpha)$ had produced the measured plant output/input data $(y, u) \in \mathbf{S}^\ell \times \mathbf{S}^\ell$, then the reference input **would have been** *the sequence $r(\alpha)$ satisfying*

$$u = C(\alpha)\,(r(\alpha) - y) \tag{62}$$

Assuming the indicated inverse exists,

$$r(\alpha) = y + C(\alpha)^{-1} u \tag{63}$$

Thus, the error, $\varepsilon(\alpha)$, **would have been**,

$$\begin{aligned} \varepsilon(\alpha) &= y - T_{\text{ref}}\ r(\alpha) \\ &= S_{\text{ref}}\ y - T_{\text{ref}} C(\alpha)^{-1} u \end{aligned} \tag{64}$$

with

$$S_{\text{ref}} = 1 - T_{\text{ref}}$$

Hence, the set of all controller parameters that achieve a performance level, ρ, **would have been**,

$$\mathbf{A}_\ell(\rho) = \{\ \alpha \in \mathbf{R}^q\ \mid \varepsilon(\alpha) = \Delta r(\alpha), \quad \Delta \in \mathbf{\Delta}(\rho)\ \} \tag{65}$$

By this argument, $\mathbf{A}_\ell(\rho)$ is the the set of all controller parameters which are unfalsified by the available data, with respect to performance level ρ.

If $\Delta \in \mathbf{NTI}$, then

$$\begin{aligned}\mathbf{A}_\ell(\rho) &= \left\{ \alpha \in \mathbf{R}^q \;\middle|\; \big\|\varepsilon(\alpha)\big\|_{\mathbf{L}_2[1,t]} \le \rho \big\|r(\alpha)\big\|_{\mathbf{L}_2[1,t]} \,, \quad t \in [1,\ell] \right\} \\ &= \bigcap_{t\in[1,\ell]} \left\{ \alpha \in \mathbf{R}^q \;\middle|\; \big\|\varepsilon(\alpha)\big\|_{\mathbf{L}_2[1,t]} \le \rho \big\|r(\alpha)\big\|_{\mathbf{L}_2[1,t]} \right\}\end{aligned} \tag{66}$$

This last expression shows that since $\mathbf{A}_\ell(\rho)$ is an intersection of inequalities, one for each data point, it follows that as more data is recorded, *i.e.*, as ℓ increases, the unfalsified parameter set, $\mathbf{A}_\ell(\rho)$, can only get smaller. The *falsified* parameter set, its complement, therefor, can only increase. Similarly, if $\Delta \in \mathbf{LTI}$, then

$$\mathbf{A}_\ell(\rho) = \left\{ \alpha \in \mathbf{R}^q \;\middle|\; \mathcal{T}\left\{\varepsilon(\alpha)\right\}^T \mathcal{T}\left\{\varepsilon(\alpha)\right\} - \rho^2 \mathcal{T}\left\{r(\alpha)\right\}^T \mathcal{T}\left\{r(\alpha)\right\} \le 0 \right\} \tag{67}$$

In this case as ℓ increases, the Toeplitz matrices get larger thereby incorporating more constraints imposed by the data.

Having obtained the unfalsified set of controller parameters, essentially a set of candidate controllers, it remains to choose a controller parameter to implement in the next iteration, $\ell + 1$. An aggressive choice is one which produces the smallest ρ, *i.e.*,

$$(\alpha^\ell_{unf}, \rho^\ell_{unf}) = \arg\min \{ \; \rho \; | \; \alpha \in \mathbf{A}_\ell(\rho) \; \} \tag{68}$$

A cautious choice, reflecting the distribution of elements in $\mathbf{A}_\ell$, is the average, or the geometric center of the set, *i.e.*,

$$(\alpha^\ell_{unf}, \rho^\ell_{unf}) = \arg avg \{ \; \rho \; | \; \alpha \in \mathbf{A}_\ell(\rho) \; \} \tag{69}$$

No matter the choice, the control parameters are considered likely to be updated whenever the unfalsified performance level, ρ^ℓ_{unf}, is smaller than the measured performance, ρ^ℓ_{meas}. Thus, the controller parameter update rule is,

$$\alpha^{\ell+1} = \begin{cases} \alpha^\ell_{unf}, \; \rho^\ell_{unf} < \rho^\ell_{meas} \\ \alpha^\ell, \quad \rho^\ell_{unf} \ge \rho^\ell_{meas} \end{cases} \tag{70}$$

Clearly if the exogenous inputs are not sufficiently rich from iteration to iteration, then it is likely that the control will not switch.

As discussed in [34], there are several advantages to this data-based control design approach:

1. The approach is nonconservative; *i.e.*, it gives "if and only if" conditions on the candidate controller $C(\alpha)$ to be unfalsified.

2. The unfalsified set of candidate controllers is determined from past data only – *no candidate controller is implemented if it is falsified*, and the test is applied *without actually implementing the candidate controller* $C(\alpha)$. Equally important, if the test fails, those candidate controllers, $C(\alpha)$, which have been falsified, again without implementation, *can all be discarded from any future consideration*.
3. The test for controller unfalsification is "plant-model free." No plant model is needed to test its conditions. It depends only on the data, the controller, and the specification.
4. The data which falsifies a controller may be open loop data or data generated by some other control law which may or may not be in the parametric set.
5. Controller falsification implies falsification of any underlying uncertainty model for the plant model, based on the same data, which would have resulted in the same controller. The converse, however, is not true: a falsified uncertainty model of the plant does not imply falsification of a controller based on this falsified uncertainty model. As a result, using the same data set, direct controller unfalsification can produce less conservative control than plant unfalsification followed by robust control design.

5.1 Relation to Windsurfing

A key element in the windsurfing approach to adaptation is to vary the reference system bandwidth. This is one of many possible performance parameters, *e.g.*, overshoot, settling time, and so on. To illustrate the idea let β denote the reference system bandwidth and let $T_{\rm ref}(\beta)$ denote the reference system. Now the set of unfalsified control parameters depends not only on ρ, the performance error, but also on β,

$$\begin{aligned} \mathbf{A}_\ell(\rho,\beta) &= \{\ \alpha \in \mathbf{R}^q \ \mid \varepsilon(\alpha,\beta) = \Delta r(\alpha), \quad \Delta \in \mathbf{\Delta}(\rho)\ \} \\ \varepsilon(\alpha,\beta) &= y - T_{\rm ref}(\beta)\ r(\alpha) \end{aligned} \tag{71}$$

We can now generate a *family* of unfalsified parameter updates which depend on β. For example,

$$(\alpha^\ell_{unf}(\beta), \rho^\ell_{unf}(\beta)) = \arg\min \{\ \rho \ \mid \alpha \in \mathbf{A}_\ell(\rho,\beta)\ \} \tag{72}$$

This will be illustrated in the example in section 6.1.

Computational Issues

Measured performance, ρ_{meas}, can be obtained directly from either (60) or (61). However, solving the the optimization problem for $(\rho_{unf}, \alpha_{unf})$ has two difficulties. First, it is not in general convex, hence, there is no guaranty of finding the optimum. For PID and/or lead-lag type controllers, which have

a small number of parameters, a combinatorial search is very effective as has been shown in [34] and will be demonstrated in the example in section 6.1.

Difficulties with the optimization are to be expected, because in essence, we are trying to solve the fixed-order control design problem, which is generically hard even when the plant is known, *e.g.*, [15]. In the case here, the plant is not known, and the problem is compounded further by using data! However, as in the output error identification problem, which is related, there are some instances where there are no local minima for parameters restricted to a region where a certain transfer function is passive [27, Ch.10,p.301]. Even if this could be applied here, obviously more assumptions about the plant are required.

The second issue is that the problem size increases as time goes on because more data is recorded, effectively adding more constraints. We offer an approximate solution to this problem in the next section on iterative unfalsification. Essentially the data is only recorded over a fixed length window which slides along with current time. This makes sense if the data record length, ℓ, is long compared to the assumed system memory.

Dealing with the ever increasing problem size can be addressed by developing recursive methods which provide outer and inner ellipsoid bounds on earlier data, hence compressing the earlier data into matrices on the order of the parameters. This is akin to least-squares estimation which compresses prior information into parameter sized covariance matrices. (This is still an idea.)

For LTI closed-loop uncertainty sets, the unfalsification tests can be recast as convex optimizations. In [37,38] it is shown that a reformulation of the performance specification as a receding horizon criterion allows the problem to be cast as a convex optimization. Another approach is described next.

5.2 LTI Unfalsification

Consider the closed-loop specification represented by the uncertainty model set (59),

$$\mathbf{J}(\rho) = \left\{ \; y, r \;\middle|\; y - T_{\rm ref}\; r = \Delta r, \quad \|\Delta\|_{\mathbf{H}_\infty} \le \rho \; \right\} \tag{73}$$

To test the candidate controller,

$$u = C(\alpha)(r - y)$$

form

$$r(\alpha) = y + C(\alpha)^{-1} u$$

where (y, u) are given data of length ℓ. The closed-loop uncertainty model is then

$$(1 - T_{\rm ref})y - T_{\rm ref} C(\alpha)^{-1} u = \Delta(y + C(\alpha)^{-1} u), \quad \|\Delta\|_{\mathbf{H}_\infty} \le \rho$$

Suppose that $C(\alpha)$ is parametrized as given by (21). If ℓ is sufficiently large, then both sides of the the above LTI uncertainty model can be operated on by $N(\alpha)$ to give,

$$\begin{aligned} e(\alpha) &= \Delta f(\alpha), \quad \|\Delta\|_{\mathbf{H}_\infty} \leq \rho \\ e(\alpha) &= N(\alpha)(1 - T_{\text{ref}})y - D(\alpha)T_{\text{ref}}u \\ f(\alpha) &= N(\alpha)y + D(\alpha)u \end{aligned} \tag{74}$$

The set of unfalsified controller parameters is then,

$$\mathbf{A}_\ell(\rho) = \left\{ \alpha \in \mathbf{R}^q \;\middle|\; \mathcal{T}\left\{e(\alpha)\right\}^T \mathcal{T}\left\{e(\alpha)\right\} - \rho^2 \mathcal{T}\left\{f(\alpha)\right\}^T \mathcal{T}\left\{f(\alpha)\right\} \leq 0 \right\} \tag{75}$$

Since $N(\alpha)$ and $D(\alpha)$ are affine in α, so are $e(\alpha)$ and $f(\alpha)$. Hence, for a fixed value of ρ, the above set is convex in α. Solving for α which produces the smallest value of ρ is in the form of a "generalized eigenvalue problem" and can be solved using convex programming [6].

6 Iterative Controller Unfalsification

In this section the unfalsification paradigm is used to develop an iterative direct adaptive controller.

During the i-th iteration (data collection period) the controller is held fixed at $C(\alpha^i) \in \mathbf{LTI}$, *i.e.*,

$$u = C(\alpha^i)(r - y) \tag{76}$$

Suppose each data collection period contains ℓ samples, where $(y^i, u^i, r^i) \in \mathbf{S}^\ell \times \mathbf{S}^\ell \times \mathbf{S}^\ell$ is the data measured. Define the corresponding *measured* performance error by,

$$\rho^i_{meas} = \min \left\{ \rho \;\middle|\; \varepsilon^i = \Delta r^i, \quad \Delta \in \mathbf{\Delta_{NTI}}(\rho) \right\} \tag{77}$$

where

$$\varepsilon^i = y^i - T_{\text{ref}}\ r^i \tag{78}$$

> **Caveat emptor** – *This definition of measured performance for the i-th data collection period is reasonable only if the period data length, ℓ, is sufficiently large so as to make negligible any effects due to controller adjustments or exogenous disturbances in previous periods.* **Assume from now on that this is the case.**

From (51) it follows that the measured performance is given by,

$$\rho^i_{meas} = \min \left\{ \rho \;\middle|\; \left\|\varepsilon^i\right\|_{\mathbf{L}_2[1,t]} \leq \rho \left\|r^i\right\|_{\mathbf{L}_2[1,t]}, \quad t \in [1, \ell] \right\} \tag{79}$$

Based solely on the data collected in the i-th period, the set of unfalsified controller parameters that achieve a performance level, ρ, is,

$$\mathbf{A}_\ell^i(\rho) = \left\{ \alpha \in \mathbf{R}^q \;\middle|\; \left\|\varepsilon^i(\alpha)\right\|_{\mathbf{L}_2[1,t]} \leq \rho \left\|r^i(\alpha)\right\|_{\mathbf{L}_2[1,t]} , \quad t \in [1,\ell] \right\} \tag{80}$$

It follows that the set of all controller parameters which are unfalsified, with respect to performance level ρ, up to and including the i-th interval, is the intersection of these sets, *i.e.*,

$$\mathbf{A}^i(\rho) = \bigcap_{k \in [1,i]} \mathbf{A}_\ell^k(\rho) \tag{81}$$

As before, it remains to choose a controller to implement in the next iteration. The aggressive choice produces the smallest ρ, *i.e.*,

$$(\alpha_{unf}^i, \rho_{unf}^i) = \arg\min \left\{ \rho \;\middle|\; \alpha \in \mathbf{A}^i(\rho) \right\} \tag{82}$$

whereas the cautious choice, reflecting the distribution of elements in $\mathbf{A}^i$, is the average, or the geometric center of the set, *i.e.*,

$$(\alpha_{unf}^i, \rho_{unf}^i) = \arg avg \left\{ \rho \;\middle|\; \alpha \in \mathbf{A}^i(\rho) \right\} \tag{83}$$

We then propose to update the control parameters whenever the unfalsified performance level, ρ_{unf}^i, is smaller than the best measured performance,

$$\rho_{meas}^k = \min_{j \in [1,i]} \rho_{meas}^j \tag{84}$$

If not, then control is returned to $C(\alpha^k)$, the controller which produced the best measured performance. Thus, the controller parameter update rule is,

$$\alpha^{i+1} = \begin{cases} \alpha_{unf}^i, \; \rho_{unf}^i < \rho_{meas}^k \\ \alpha^k, \quad \rho_{unf}^i \geq \rho_{meas}^k \end{cases} \tag{85}$$

This is a slightly different procedure than in the previous "one-step-at-a-time" case. Here, because the control is held fixed at $C(\alpha^i)$ for a long time, we have (we assume) a good reading of the performance with this control. In the previous formulation, the control can switch at every instant when new data is acquired.

6.1 Simulation Example: PI Control of a Nonlinear System

The iterative procedure is simulated with the following nonlinear plant system:

$$y = G\Big(v + N(u)\Big)$$

$$G = \frac{.1z^{-1}}{1 - .4z^{-1}}$$

$$N(u) = \begin{cases} 0, & |u| \le d \\ u - d, \; u > 0 \\ u + d, \; u < 0 \end{cases} \tag{86}$$

$$\big\|v\big\|_{rms} \le \sigma$$

The plant system is thus a linear system, G, driven by an RMS-bounded disturbance, v, and controlled through a deadband nonlinearity, $N(\cdot)$, with deadband of size d. The control is given by the PI control,

$$u = C(\alpha)(r - y)$$

where

$$C(\alpha) = \alpha_P + \frac{\alpha_I z^{-1}}{1 - z^{-1}} \tag{87}$$

The reference system is

$$T_{\text{ref}} = \frac{(1-a)z^{-1}}{1 - az^{-1}}, \quad a = \exp(-2\pi f_{\text{ref}})$$

Figures 3-6 show the results of the simulations. Each figure has two rows and four columns. Each row corresponds to a different bandwidth (f_{ref}) of the reference system. The rows are as follows:

- row 1: the initial output response, before adaptation, compared to the reference system output.
- row 2: the final output response, after adaptation, compared to the reference system output.
- row 3: the per iteration values of ρ^i_{meas}, ρ^i_{unf}, and the $\mathbf{H}_\infty$-norm of the error between the linearized system and the reference system.
- row 4: the PI gains per iteration.

The simulations were performed under the following conditions:

- The control was initialized as the low gain integrator:

$$C = \frac{.01z^{-1}}{1 - z^{-1}}$$

- A single repeating cycle of the reference input is given by:

$$r = \begin{cases} 1 & t = 1:200 \\ -1 & t = 201:400 \\ 0 & t = 401:600 \end{cases}$$

- There are two cycles of 4 iterations each of this reference.
  - During cycle 1, the reference system bandwidth
    $f_{\text{ref}} = .005$ hz
    The results are shown in column 1 of all the figures.
  - During cycle 2, the reference system bandwidth
    $f_{\text{ref}} = .05$ hz
    The results are shown in column 2 of all the figures.
- The deadband width (d) and RMS-disturbance level (σ) were set as follows:

  figure 3 $d = 0 \quad \sigma = 0$
  figure 4 $d = 0 \quad \sigma = .1$
  figure 5 $d = 1.5$ $\sigma = 0$
  figure 6 $d = 1.5$ $\sigma = .1$

We see in all cases that the iterative unfalsified adaptation works very well despite some extreme variations and no prior knowledge about the plant system. Although not shown, the intermediate time responses are not very much different than the final responses (after 4 iterations).

7 Concluding Remarks

There are several intriguing aspects of unfalsification as applied to direct adaptive control. First, existing data can be used to falsify an experiment you would like to perform, but cannot. Secondly, controllers can be proven to be unable to meet the closed-loop performance specification without being implemented. This reduces the set of unfalsified controllers, and this reduction is non-conservative. But what about convergence? The answer to this could be: *why convergence?* If adaptation is meant to be used in the face of highly uncertain systems, which may exhibit large variations over time and operating conditions, there is no convergence. We just keep throwing away bad controllers. A well respected American football player, when asked why he was such a good defender against the run, replied, "I just keep knock'n 'em down 'till I get to the one with the ball."

An issue not discussed here is the general topic of estimating an uncertainty model from data. As shown in [22], unfalsification of an uncertainty model consisting of a Gaussian white disturbance and an $\mathbf{H}_\infty$ bounded model error leads to conclusions, in the control design context, that are not all that different from other methods, *e.g.*, *stochastic embedding* [10] and *model-error-modeling* [28]. All these methods are utilizing "least-squares" in some form,

and perhaps that is the connecting piece, despite the fact that the interpretations seem different. Although often extremely effective, I view all these methods as somewhat piece-meal, and not satisfactorily elegant. Perhaps this is due to fundamental limitations, *e.g.*, it is not possible to uniquely separate dynamic and disturbance uncertainty from finite data. For control design, however, this limitation may have no consequence. No model is needed in direct controller falsification [34] or direct tuning [12,14] – the data "speaks" to the closed-loop specification.

The answer will come, I believe, by resolving what I consider a fundamental challenge: *to develop a self consistent theory of identification and validation from a single set of assumptions.*

Acknowledgements

Brian D. O. Anderson.

References

1. B.D.O. Anderson, R.R. Bitmead, C.R. Johnson, Jr., P.V. Kokotovic, R.L. Kosut, I.M.Y. Mareels, L. Praly, and B.D. Riedle, (1986), *Stability of Adaptive Systems: Passivity and Averaging Analysis*, MIT Press, 1986.
2. B.D.O.Anderson, and R.L.Kosut, Adaptive Robust Control: On-Line Learning, Proc CDC'91, Brighton, England, Dec 1991.
3. K. J. Åström and B. Wittenmark, *Adaptive Control*, Addison-Wesley, 1995.
4. K. J. Åström and J. Nilsson, "Analysis of a scheme for iterated identification and control," *Proc. IFAC Symp. on Identification*, pp. 171-176, Copenhagen, Denmark, 1994.
5. K. J. Åström , "Matching criteria for control and identification," *Proc. 2nd European Control Conference*, pp.248-251, Groningen, Holland, June 1993.
6. S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*, SIAM Studies in Applied Mathematics, 1994.
7. P. B. Brugarolas, V. Fromion, and M. G. Safonov, "Robust switching missile autopilot", *Proc. 1998 ACC*, pp.3665-3669, Philadelphia, PA, June 24-26,1998.
8. C. Foias and A. E. Frazho, *The Commutant Lifting Approach to Interpolation Problems*, Birkhauser-Verlag, 1990.
9. M. Gevers, "Towards a joint design of identification and control," *Essays on Control: Perspectives in the Theory and its Applications*, editors: H.L. Trentelman and J.C. Willems, pp. 111-151, Birkhauser, Boston, MA, 1993.
10. G.C. Goodwin, M. Gevers, and B. Ninness, "Quantifying the error in estimated transfer functions with applications to model order selection," *IEEE Trans. on Automatic Control*, vol. 37, no.7, pp. 913-928, July 1992.
11. U. Grenander and G. Szebo, *Toeplitz Forms and Their Applications*, Chelsea Publishing Co., New York, 1958.
12. H. Hjalmarsson, S. Gunnarsson, and M. Gevers, "A convergent iterative restricted complexity control design scheme," *Proc. 1994 CDC*,pp. 1735-1740, Orlando, Florida, Dec. 1994.

13. H. Hjalmarsson, S. Gunnarsson, and M. Gevers, "Optimality and suboptimality of iterative identification and control design schemes," *Proc. 1995 ACC*, pp.2559-2563, Seattle, Washington, June 1995.
14. L. Kammer, R. Bitmead, P. Bartlett, "Direct iterative tuning via spectral analysis," *Proc. 1998 CDC*, pp.2874-2879, Tampa, FL, Dec. 1998.
15. R. L. Kosut, "Sub-optimal control of linear time-invariant systems subject to control structure constraints," *IEEE Trans. on Auto. Contr.*, AC-15(5):557-563, June 1970.
16. R. L. Kosut, "Adaptive calibration: An approach to uncertainty modeling and on-line robust control design," *Proc. 25th IEEE CDC*, Athens, Greece, Dec. 1986.
17. R.L. Kosut, G.C. Goodwin, and M.P. Polis, *editors*, "Special Issue on System Identification for Robust Control Design," *IEEE Trans. on Automatic Control*, vol. 37, no.7, pp. 899-1008, July 1992.
18. R. L. Kosut, M. Lau, and S. Boyd, "Set-membership identification of systems with parametric and nonparametric uncertainty," "Special Issue on Identification for Robust Control Design," *IEEE Trans. Aut. Contr.*, vol. 37, no. 7, July, 1992.
19. R.L. Kosut, "Uncertainty model unfalsification: a system identification paradigm compatible with robust control design," *Proc. 1995 CDC*, New Orleans, LA, Dec. 1995.
20. R. L. Kosut, "Iterative adaptive robust control via uncertainty model unfalsification," *Proc. 1996 IFAC World Congress*, June, 1996, San Francisco, CA
21. R. L. Kosut and B. D. O. Anderson, "Uncertainty model unfalsification," *Proc. 1997 CDC*, San Diego, CA, Dec, 1997.
22. R. L. Kosut, "Uncertainty model unfalsification," CD-ROM SYSID 2000, Santa Barbara, CA, 21-23 June 2000.
23. M. V. P. Kruger and K. Poolla, "Validation of uncertainty models in the presence of stochastic disturbance signals," *Selected Topics in Modelling and Control*, Vol. 11,pp.1-8, Delft University Press, Dec. 1998.
24. W.S. Lee, B.D.O. Anderson, R.L. Kosut, and I.M.Y. Mareels, "A new approach to adaptive robust control," *Int. J. of Adaptive Control and Signal Processing*, vol. 7, pp. 183-211, 1993.
25. W.S. Lee, B.D.O. Anderson, I.M.Y. Mareels, and R.L. Kosut, "On some key issues in the windsurfer approach to robust adaptive control", *Automatica*, Vol. 31, No. 11, pp. 1619-1636, 1995.
26. M.M. Livestone, M.A. Dahleh, and J.A. Farrell, "A framework for robust control based model invalidation", *Proc. 1994 ACC*, pp. 3017-3020, Baltimore, MD, June 1994.
27. L. Ljung, *System Identification: Theory for the User*, First Edition, Prentice-Hall, 1987.
28. L. Ljung, "Model validation and model error modeling," *The Åström Symposium on Control*, pp.15-42, Lund, Sweden, August 27-28, 1999.
29. M. Massoumnia and R. L. Kosut, "A family of norms for system identification problems," *Proc. 1993 ACC*, San Francisco, CA, June 1993, and *IEEE Trans. Aut. Control*, vol. 39, no. 5, pp. 1027-1031, May 1994.
30. M. P. Newlin and R. S. Smith, "A generalization of the structured singular value and its application to model validation," *IEEE Trans. Aut. Contr.*, vol. 43, no. 7, pp. 901-907, July 1998.

31. F. Paganini, "A set-based approach for white noise modeling," *IEEE Trans. Aut. Contr.*, vol. 41, no. 10, pp. 1453-1465, Oct. 1996.
32. S. M. Phillips, R. L. Kosut, G. F. Franklin, "An averaging analysis of discrete-time indirect adaptive control", *Proc. 1988 American Control Conference*, Atlanta, GA, 1988.
33. K. Poolla, P. Khargonnekar, A. Tikku, J. Krause, and K. Nagpal, "A time-domain approach to model validation", *Proc. 1992 ACC*, Chicago, IL, June 1992, and *IEEE Trans. Aut. Contr.*, vol. 39, no. 5, pp. 951-959, May 1994.
34. M. G. Safonov and T. C. Tsao, "The unfalsified control concept and learning", *IEEE Trans. Aut. Contr.*, vol. 42, no. 6, pp. 843-847, June 1997.
35. R. S. Smith and J. C. Doyle, "Model validation – a connection between robust control and identification," *Proc. 1989 ACC*, pp. 1435-1440, June 1989, and *IEEE Trans. Aut. Contr.*, vol. 37, pp. 942-952, July 1992.
36. B. Wahlberg and P.M. Makila (1996), "On approximation of stable linear dynamical systmes using Laguerre and Kautz functions," *Automatica*, 32(5), 693-708.
37. B. R. Woodley, R. L. Kosut, and J. P. How, "Uncertainty model unfalsification with simulation," *Proc. 1998 ACC*, pp. 2754-2755, Philadelphia, PA, June 24-26, 1998.
38. B. R. Woodley, J. P. How, and R. L. Kosut, "Direct unfalsified controller design - solution via convex optimization," submitted 1999 ACC, San Diego, CA, June 2-4, 1999.
39. Z. Zang, R. R. Bitmead, and M. Gevers, "$\mathbf{H}_2$ iterative model refinement and control robustness enhancement," *Proc. 30th IEEE CDC*, pp. 279-284, Brighton, UK, Dec. 1991.
40. K. Zhou and J. C. Doyle, *Essentials of Robust Control*, Prentice Hall, 1998.

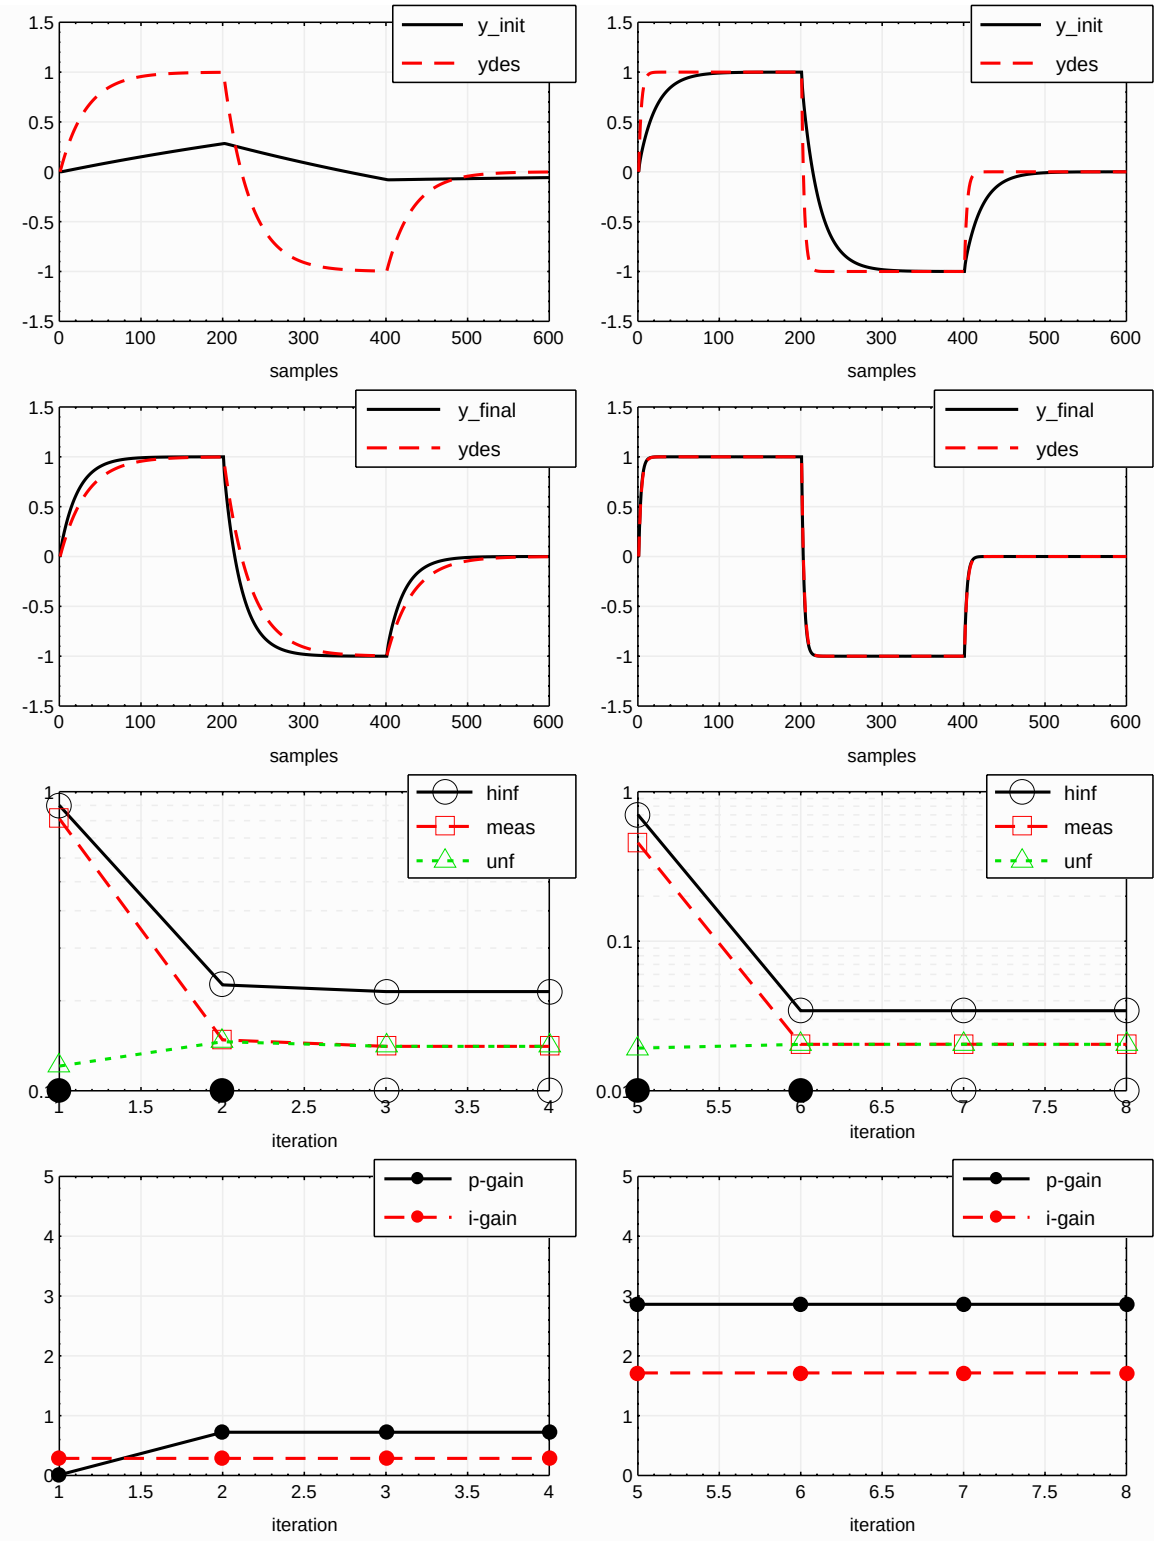

Fig. 3. Deadband width $d = 0$; RMS-disturbance $\sigma = 0$; reference bandwidth $f_{\text{ref}} = .005$ hz (col 1), $f_{\text{ref}} = .05$ hz (col 2).

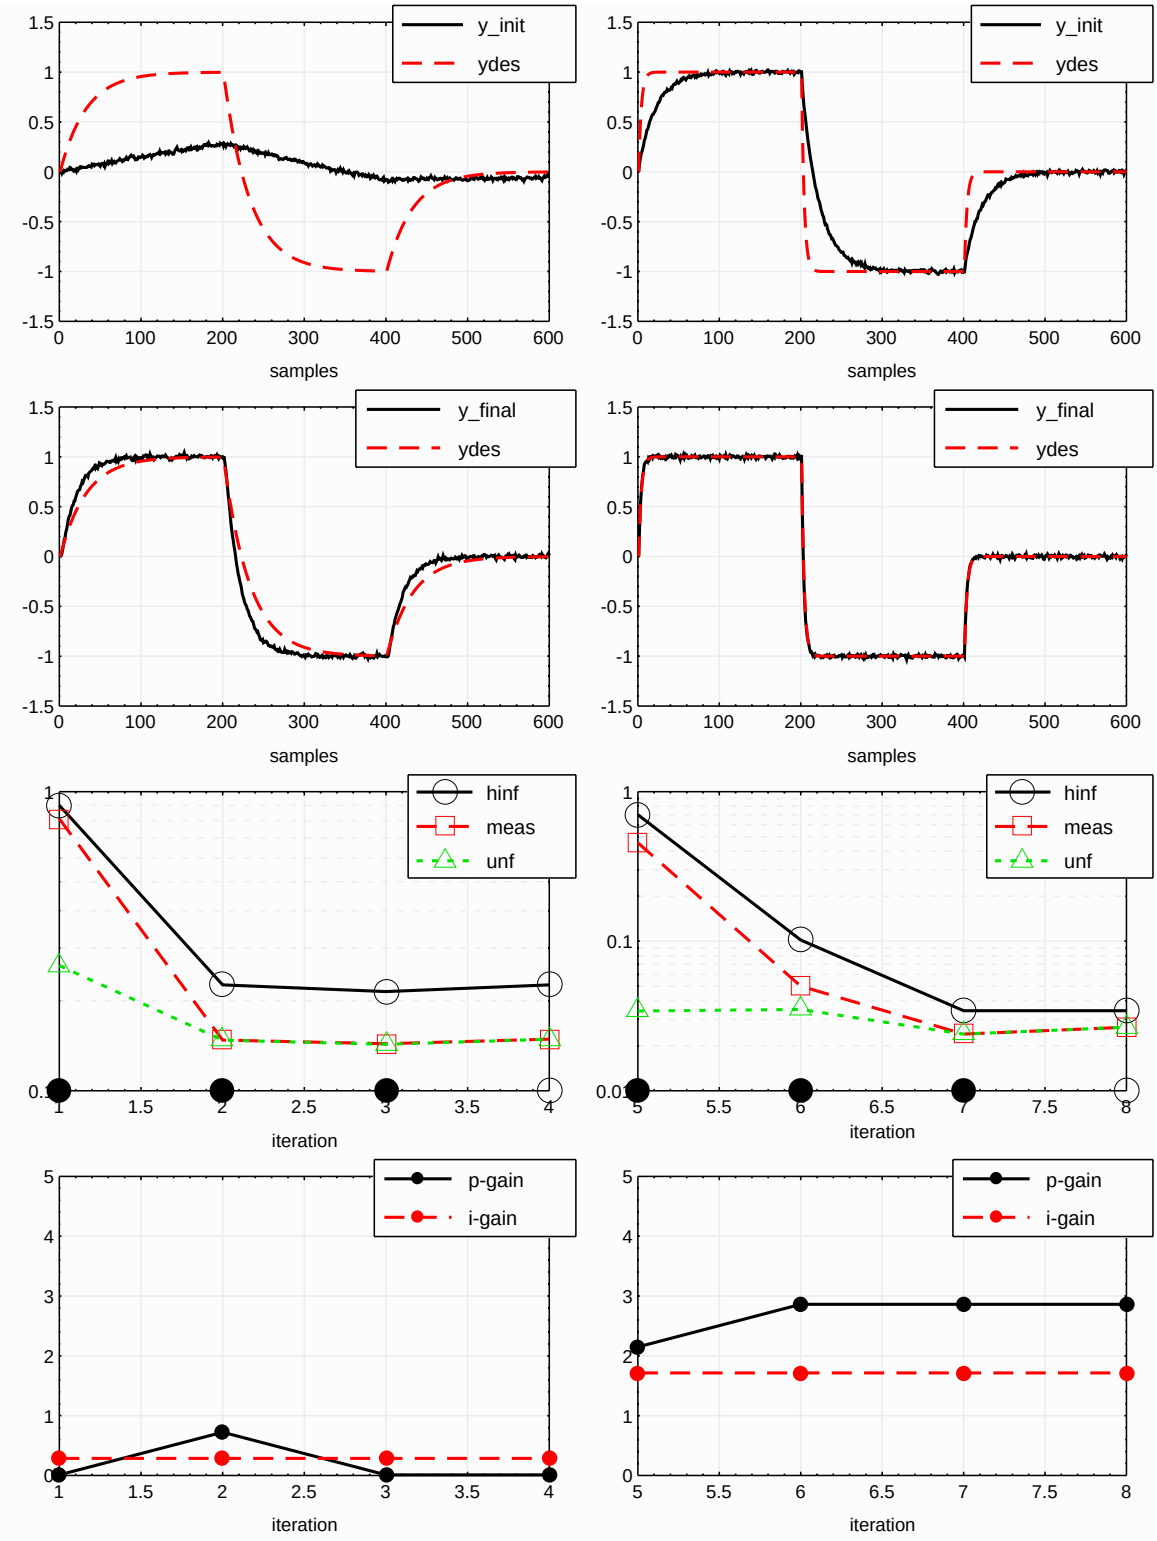

Fig. 4. Deadband width $d = 0$; RMS-disturbance $\sigma = .1$; reference bandwidth $f_{\text{ref}} = .005$ hz (col 1), $f_{\text{ref}} = .05$ hz (col 2).

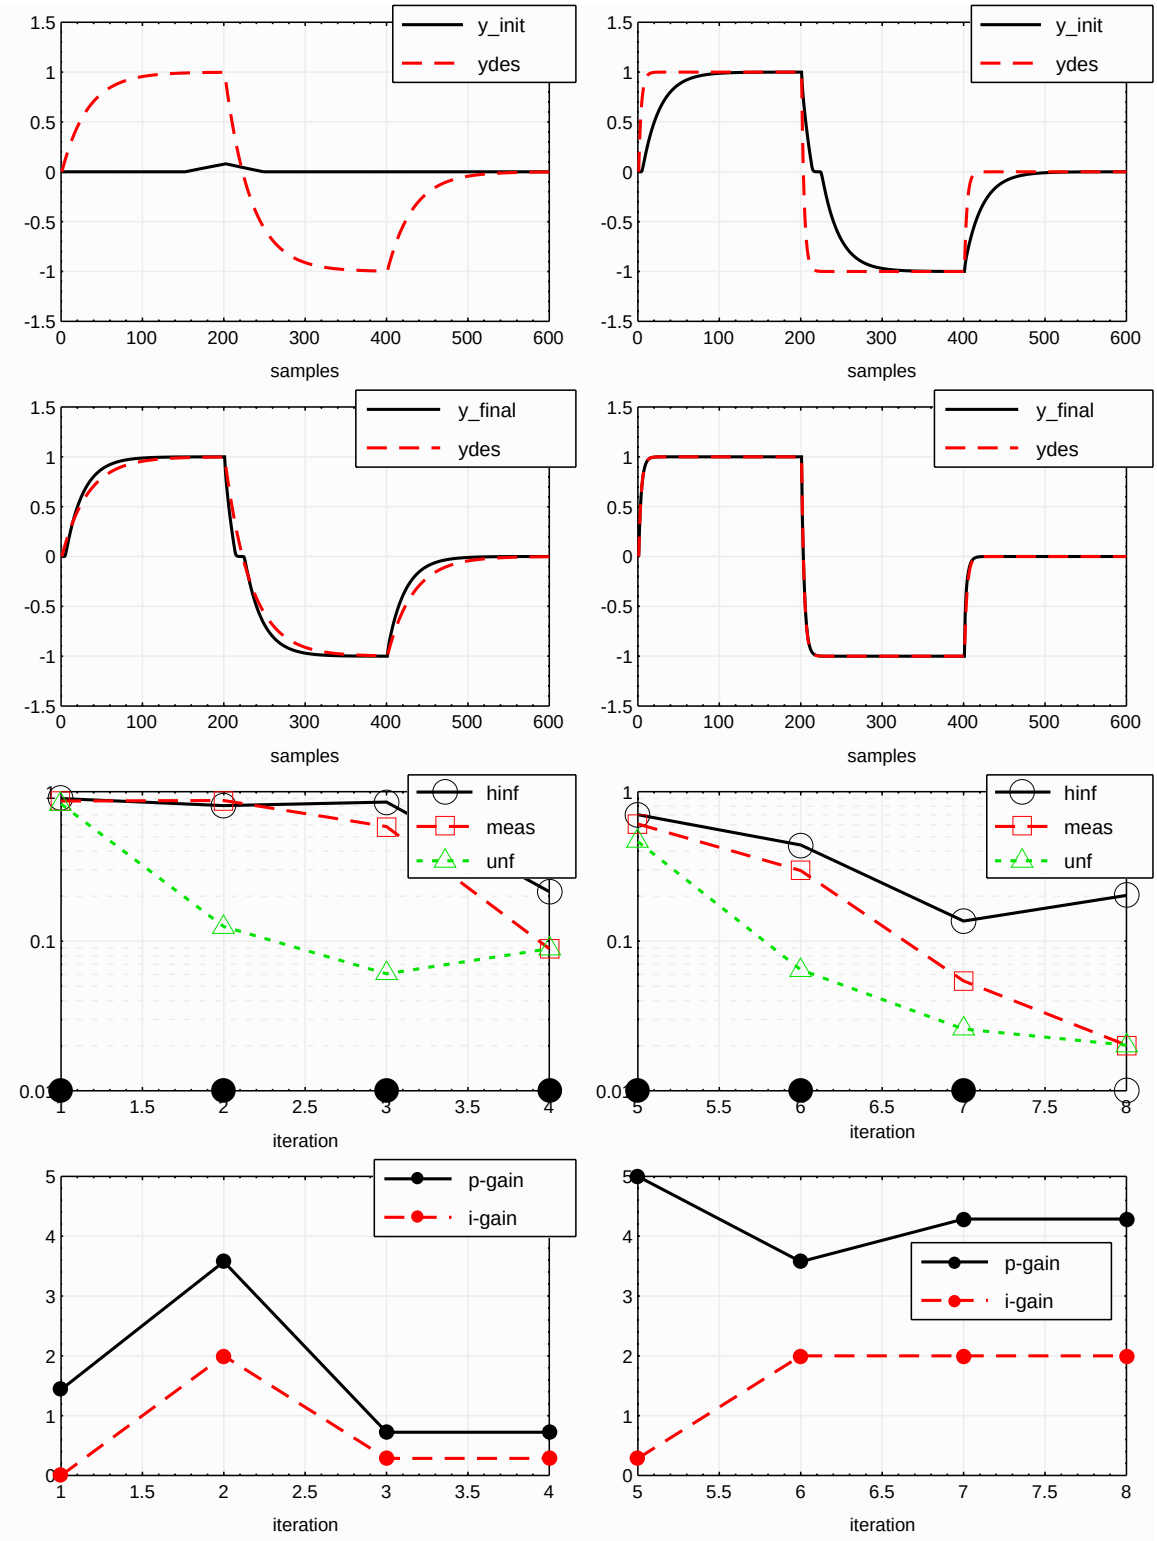

Fig. 5. Deadband width $d = 1.5$; RMS-disturbance $\sigma = 0$; reference bandwidth $f_{\text{ref}} = .005$ hz (col 1), $f_{\text{ref}} = .05$ hz (col 2).

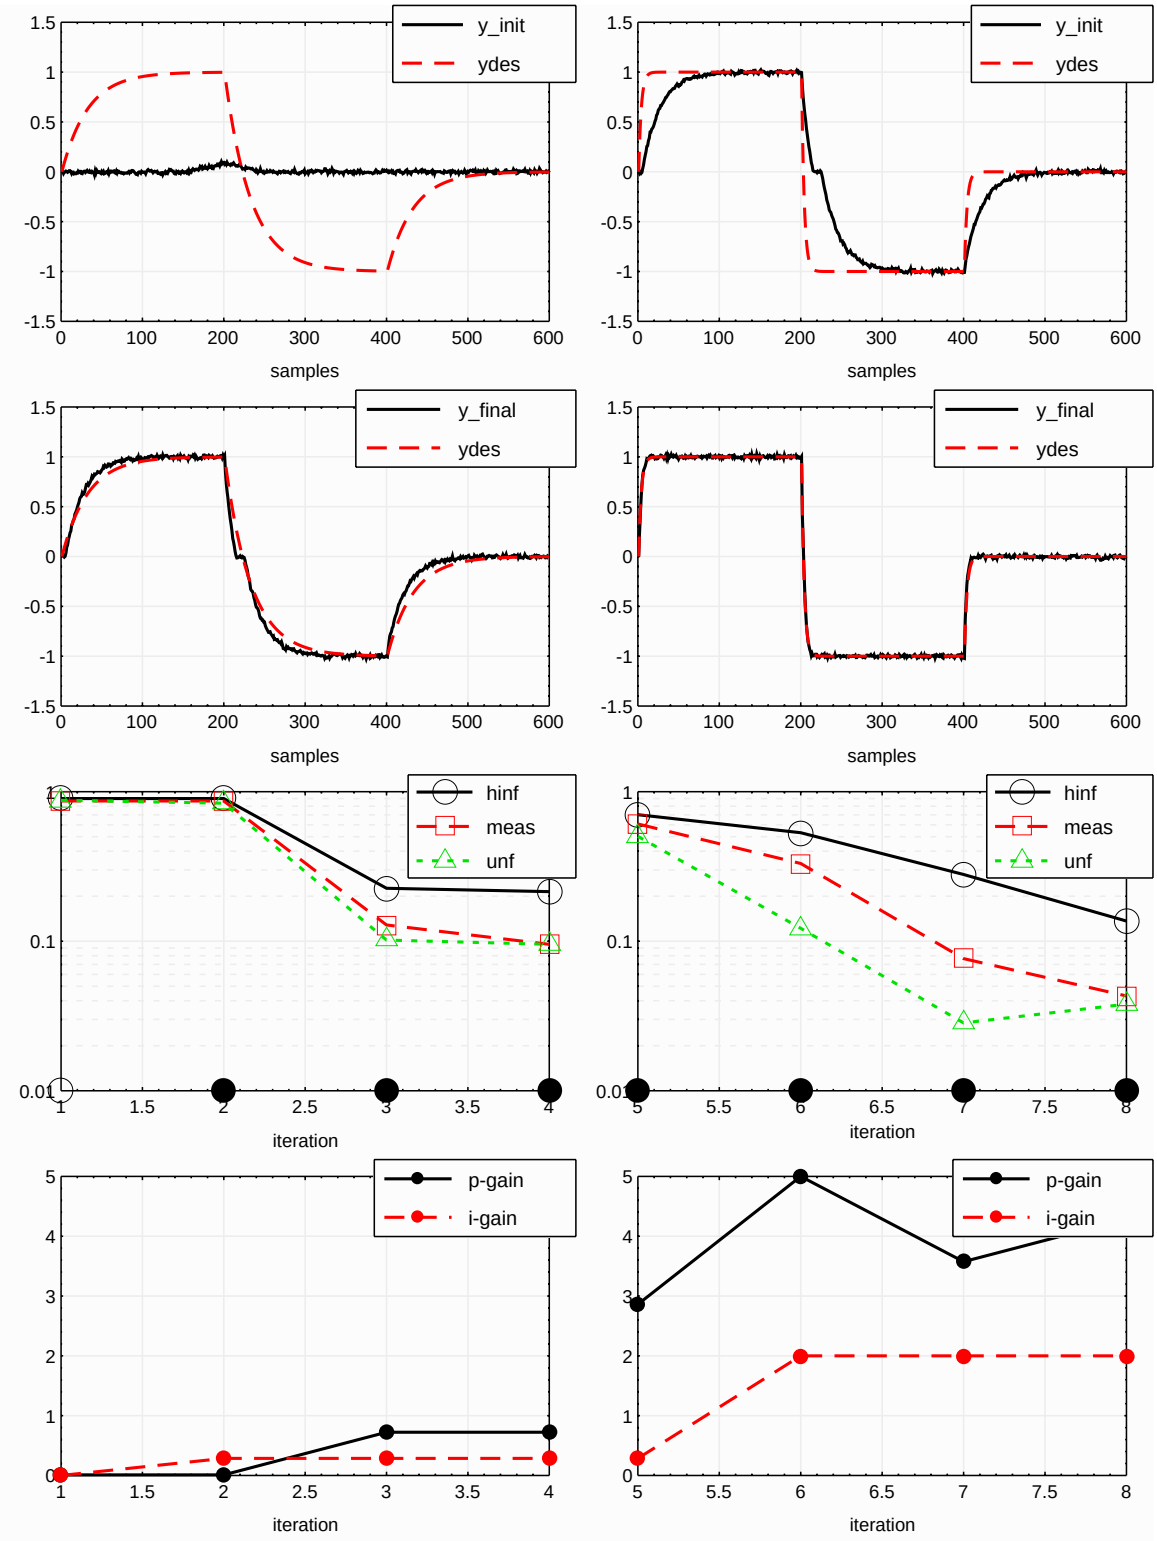

Fig. 6. Deadband width $d = 1.5$; RMS-disturbance $\sigma = .1$; reference bandwidth $f_{\text{ref}} = .005$ hz (col 1), $f_{\text{ref}} = .05$ hz (col 2).