How many experiments are needed to adapt?*

Sergio Bittanti[1], Marco C. Campi[2], and Maria Prandini[1]

[1] Dipartimento di Elettronica e Informazione - Politecnico di Milano, piazza Leonardo da Vinci 32, 20133 Milano, Italia. E-mail: {bittanti,prandini}@elet.polimi.it
[2] Dipartimento di Elettronica per l'Automazione - Università di Brescia, via Branze 38, 25123 Brescia, Italia. E-mail: marco.campi@ing.unibs.it

Summary. System design in presence of uncertainty calls for experimentation, and a question that arises naturally is: how many experiments are needed to come up with a system meeting certain performance requirements?
This contribution represents an attempt to answer this fundamental question. Results are confined to a specific set-up where adaptation is performed according to a worst-case perspective, but many considerations and reflections are central to adaptation in general.

1 Introduction

Given a system $\mathcal{S}$, consider the problem of designing a device $\mathcal{D}$ that achieves some desired behavior when interacting with $\mathcal{S}$. The specification of the 'desired behavior' depends on the intended use of the device, and is usually expressed in terms of some signal $s_{\mathcal{D}}(\omega)$, with reference to certain operating conditions $\omega \in \Omega$ of interest (Figure 1).

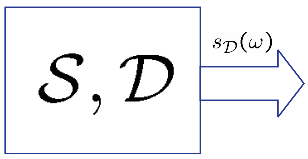

Fig. 1. Characterization through signal $s_{\mathcal{D}}(\omega)$ of device $\mathcal{D}$ while interacting with system $\mathcal{S}$ in the operating condition ω

* This work is supported by MIUR (Ministero dell'Istruzione, dell'Università e della Ricerca) under the project *Identification and adaptive control of industrial systems* and by CNR - IEIIT.

Example 1 (simulator). Suppose that the device should act as a simulator of the system when the system input u takes on value in a given class of signals U. In this case, $\omega = u$ and the desired behavior for the device can be expressed in terms of the multidimensional signal $s_{\mathcal{D}}(u) = (y(u), y_{\mathcal{D}}(u))$, where $y_{\mathcal{D}}(u)$ and $y(u)$ represent the outputs of the device and of the system fed by the same input signal $u \in U$ (see Figure 2). Signal $s_{\mathcal{D}}(u)$ should be such that $y_{\mathcal{D}}(u) \simeq y(u)$, for every operating condition of interest, that is for every $u \in U$. $\square$

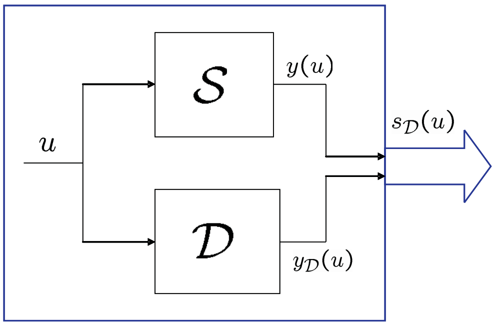

Fig. 2. Device $\mathcal{D}$ acting as a simulator of system $\mathcal{S}$

Example 2 (disturbance compensator). Suppose that the output of system $\mathcal{S}$ is affected by some additive disturbance and the device $\mathcal{D}$ is introduced for compensating the disturbance according to the feedforward scheme in Figure 3. In this case the operating condition is defined by the disturbance realization d. If we denote by $y_{\mathcal{D}}(d)$ the controlled output of the system when the disturbance realization is d, then the desired behavior can be expressed in terms of the signal $s_{\mathcal{D}}(d) = y_{\mathcal{D}}(d)$ and $s_{\mathcal{D}}(d)$ should be small for every d in some set D. $\square$

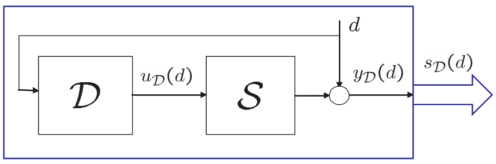

Fig. 3. Device $\mathcal{D}$ acting as a disturbance compensator for system $\mathcal{S}$

Devising a suitable $\mathcal{D}$ for a system $\mathcal{S}$ requires knowledge of some sort on $\mathcal{S}$. Most literature in science and engineering relies on a model-based approach, namely it is assumed that a mathematical model for $\mathcal{S}$ is a-priori available. Alternatively, the knowledge on $\mathcal{S}$ can be accrued through experimentation. This latter approach, considered herein, is referred to as 'adaptive design', [1, 2, 3, 4, 5], since the problem is to adapt $\mathcal{D}$ on the basis of experiments in the face of the lack of a-priori knowledge on system $\mathcal{S}$.

In adaptive design, one fundamental question to ask is:

> *How extensively do we need to experiment in order to come up with a device meeting certain performance requirements?*

This fundamental –and yet largely unanswered– question is the theme this contribution is centered around.

In this paper, a *worst-case perspective* with respect to the possible operating conditions is adopted, and we provide an answer to the above question in this specific set-up. For one answer, many more are the answers that this contribution is incapable to provide, which will also be enlightened along our way.

2 Worst-case approach to adaptation

Worst-case performance

Suppose that the performance of device $\mathcal{D}$ operating in condition ω is quantified by a cost $c(s_{\mathcal{D}}(\omega))$. Then, the worst-case performance achieved by $\mathcal{D}$ over the set Ω of operating conditions is

$$\max_{\omega \in \Omega} c(s_{\mathcal{D}}(\omega)),$$

and, correspondingly, one wants to design

$$\mathcal{D}^\star = \arg\min_{\mathcal{D}} \max_{\omega \in \Omega} c(s_{\mathcal{D}}(\omega)). \tag{1}$$

$c^\star$ denotes the worst-case performance of device $\mathcal{D}^\star$, that is $c^\star = \max_{\omega \in \Omega} c(s_{\mathcal{D}^\star}(\omega))$.

In e.g. the simulator Example 1, $\omega = u$ and one can take $c(s_{\mathcal{D}}(u)) = \|y(u) - y_{\mathcal{D}}(u)\|_2$, the 2-norm of the error signal $y(u) - y_{\mathcal{D}}(u)$. $c^\star$ can then be interpreted as an upper bound to the largest 2-norm discrepancy between the system behavior and the behavior of the simulator $\mathcal{D}^\star$ in the same operating condition:

$$\|y(u) - y_{\mathcal{D}^\star}(u)\|_2 \leq c^\star, \ \forall u \in U.$$

In the disturbance compensator Example 2, a sensible cost is the 2-norm $c(s_{\mathcal{D}}(d)) = \|y_{\mathcal{D}}(d)\|_2$. Then, the best disturbance compensator $\mathcal{D}^\star$ satisfies:

$$\|y_{\mathcal{D}^\star}(d)\|_2 \leq c^\star, \ \forall d \in D.$$

In many cases, the device $\mathcal{D}$ is parameterized by a vector $\gamma \in \Re^k$, in which case we write $\mathcal{D}_\gamma$ to indicate device $\mathcal{D}$ with parameter γ, and hence designing a device corresponds to selecting a value for γ. Then, with the shorthand

$$J_\gamma(\omega) := c(s_{\mathcal{D}_\gamma}(\omega)),$$

the min-max optimization problem (1) can be rewritten as the following robust optimization program with $k+1$ optimization variables:

$$\text{RP}: \min_{\gamma, c \in \Re^{k+1}} c \quad \text{subject to:} \tag{2}$$
$$J_\gamma(\omega) \leq c, \ \forall \omega \in \Omega.$$

Note that, given a γ, the slack variable c represents an upper bound on the cost $J_\gamma(\cdot)$ achieved over Ω by the device with parameter γ. By solving (2) we seek that $\gamma^\star$ that corresponds to the smallest upper bound $c^\star$.

Adaptive design

In model-based design, the cost $J_\gamma(\omega)$ can be evaluated based on the model, and then $\gamma^\star$ is found by solving the robust optimization program (2). Instead, when system $\mathcal{S}$ is unknown or only partially known, the cost $J_\gamma(\omega)$ cannot be explicitly computed so that the constraints in (2) are not known. However, one can conceive of evaluating the constraints experimentally. What exactly this means is discussed in the sequel.

Each constraint is associated with an operating condition $\omega \in \Omega$. To evaluate experimentally a constraint in a specific condition, say $\hat{\omega} \in \Omega$, that is to determine experimentally the domain of feasibility in the (γ, c)'s space where the constraint $J_\gamma(\hat{\omega}) \leq c$ holds, one should run a set of experiments, all in the $\hat{\omega}$ condition, each of which performed with a different device $\mathcal{D}_\gamma$, $\gamma \in \Re^k$. In this way $s_{\mathcal{D}_\gamma}(\hat{\omega})$ is measured for every γ and $J_\gamma(\hat{\omega})$ can be computed. An objection to this way of proceeding is that it would require in principle to test the performance achieved with every and each device $\mathcal{D}_\gamma$ in place. It is an interesting fact that in many situations the overwhelming experimental effort involved in testing many times with different $\mathcal{D}_\gamma$'s can be avoided, and just one single experiment is enough for the purpose of computing $J_\gamma(\hat{\omega})$.

Take e.g. the simulator Example 1. In this example, if $\hat{u}$ is injected into $\mathcal{S}$, signal $\hat{y} = \mathcal{S}[\hat{u}]$ can be collected, along with signal $\hat{u}$ itself. Based on this single experiment, one can then compute $y(\hat{u}) - y_{\mathcal{D}_\gamma}(\hat{u}) = \hat{y} - \mathcal{D}_\gamma[\hat{u}]$ for all γ's, where $\mathcal{D}_\gamma[\hat{u}]$ is obtained by filtering $\hat{u}$ with $\mathcal{D}_\gamma$, an operation that can be executed as an off-line post-process of signal $\hat{u}$. After $y(\hat{u}) - y_{\mathcal{D}_\gamma}(\hat{u})$ has been computed, the constraint $\|y(\hat{u}) - y_{\mathcal{D}_\gamma}(\hat{u})\|_2 = J_\gamma(\hat{u}) \leq c$ is evaluated.

The same conclusion that one experiment is enough can also be drawn for Example 2 whenever both the system and the device are linear. Indeed, swapping the order of $\mathcal{S}$ and $\mathcal{D}_\gamma$, we have:

$$y_{\mathcal{D}_\gamma}(d) = \mathcal{S}[\mathcal{D}_\gamma[d]] + d = \mathcal{D}_\gamma[\mathcal{S}[d]] + d. \quad (3)$$

If we run an experiment in which disturbance $\hat{d}$ is measured and this disturbance is also injected as input to the system (i.e. $\mathcal{D}$ is set to 1 during experimentation in the scheme of Figure 3), from the measured system output $\hat{y} = \mathcal{S}[\hat{d}] + \hat{d}$ and from $\hat{d}$ itself we can then determine

$$\begin{aligned} y_{\mathcal{D}_\gamma}(\hat{d}) &= \mathcal{D}_\gamma[\mathcal{S}[\hat{d}]] + \hat{d} \quad \text{(using (3))} \\ &= \mathcal{D}_\gamma[\hat{y} - \hat{d}] + \hat{d}, \end{aligned}$$

where computation of $\mathcal{D}_\gamma[\hat{y} - \hat{d}]$ is executed off-line similarly to the simulator example. By computing $\|y_{\mathcal{D}_\gamma}(\hat{d})\|_2 = J_\gamma(\hat{d})$ constraint $J_\gamma(\hat{d}) \leq c$ is then evaluated.

In the sequel we shall assume that one single experiment in condition $\hat{\omega}$ suffices to determine constraint $J_\gamma(\hat{\omega}) \leq c$. This assumption is not fulfilled in all applications of the adaptive scheme, and further discussion on this point is provided in Section 5.

Remark 1. The reader may have noticed that lack of knowledge, for which adaptation is required, can enter the problem in different ways. In Example 1, it was system $\mathcal{S}$ to be unknown. In the disturbance compensator Example 2, again uncertainty stayed with the system $\mathcal{S}$, but even the set D for d could be unknown.
The seemingly different nature of the uncertainty in $\mathcal{S}$ and in D can be leveled off by adopting a more abstract behavioral perspective, [6], where the system is just seen as a set of behaviors, i.e. of possible realizations of system signals. In such framework, uncertainty simply corresponds to say that the set of behaviors defining the system is not a-priori known. □

We are now facing the central issue this contribution is centered around, that is: an exact solution of the robust optimization program (2) requires to consider as many experiments as the number of elements in Ω, normally an *infinite* number. The impossibility to carry out this task suggests introducing approximate schemes where only a *finite* number of ω's, that is a *finite* number of experiments, are considered. Thus, we can at this point more precisely spell out the question we posed at the end of Section 1, and ask:

> *How many experiments do we need to perform to come up with a design that approximates the solution $\mathcal{D}^\star$ of (2) to a desired level of accuracy?*

3 The experimental effort needed for adaptation

The fact that one concentrates on a finite number of operating conditions only may appear naive. The interesting fact is that this way of proceeding can be cast within a solid mathematical theory providing us with guarantees on the level of accuracy obtained.

Fix an integer N, and let $\omega^{(1)}$, $\omega^{(2)}$, ..., $\omega^{(N)} \in \Omega$ be the operating conditions of N experiments run on the system to evaluate the N corresponding constraints for the robust program (2). The robust optimization problem restricted to the N experienced scenarios $\omega^{(i)}$, $i = 1, 2, \ldots, N$, reduces to the following finite optimization problem referred in the sequel to as 'scenario program':

$$\mathrm{SP}_N : \min_{\gamma, c \in \Re^{k+1}} c \quad \text{subject to:} \qquad (4)$$

$$J_\gamma(\omega^{(i)}) \leq c, \ i = 1, 2, \ldots, N.$$

As for the selection of the scenarios $\omega^{(i)}$, $i = 1, 2, \ldots, N$, we suppose that they are extracted from set Ω according to some probability distribution P that reflects the likelihood of the different ω situations. This is naturally the case in the disturbance compensator Example 2, assuming the environment randomly selects the disturbance realizations according to an invariant scheme. If the scenarios are selected by the designer of the experiment, like u in Example 1, probability P is artificially introduced to describe the likelihood of the different operating conditions.

Let $(\gamma_N^\star, c_N^\star)$ be the solution of SP_N. $c_N^\star$ quantifies the performance of the device with parameter $\gamma_N^\star$ over the extracted operating conditions $\omega^{(1)}$, $\omega^{(2)}$, ..., $\omega^{(N)}$. Moreover, we clearly have $c_N^\star \leq c^\star$, the optimal cost with all the constraints in place, that is, for the extracted scenarios, we have designed a very efficient device, in actual effects one that even outperforms device $\mathcal{D}^\star$. We cannot be satisfied with this sole result, however, since, due to the limited number of scenarios, there is no guarantee whatsoever with respect to the much larger multitude of possible operating conditions, all those that have not been seen when performing the design of $\gamma_N^\star$. Hence, the following question arises naturally: what can we claim regarding the performance of the designed device for all other operating conditions $\omega \in \Omega$, those that were not experienced while doing the design according to SP_N in (4)? Answering this question is necessary to provide accuracy guarantees and to pose the method on solid grounds.

The posed question is of the 'generalization type' in a learning-theoretic sense: we want to know how the solution $(\gamma_N^\star, c_N^\star)$ generalizes from experienced operating conditions to unexperienced ones. For ease of explanation, we shall henceforth concentrate on robust optimization problems of *convex-type*, since this case can be handled in the light of a powerful theory that has recently appeared in the literature of robust optimization, [7, 8]. The non-convex case can be dealt with along a more complicated approach and is not discussed herein.

RESULT: *Select a 'violation parameter' $\epsilon \in (0,1)$ and a 'confidence parameter' $\beta \in (0,1)$.*
If N satisfies

$$\sum_{i=0}^{k} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i} \leq \beta, \tag{5}$$

then, with probability no smaller than $1-\beta$, the solution (γ_N^*, c_N^*) to (4) satisfies all constraints of problem (2) with the exception of those corresponding to a set of operating conditions whose probability is at most ϵ.

Bound (5) can be found in [9], a contribution still in the general vein of the theoretical approach opened up in [7, 8].

Let us try to understand in detail the meaning of this result. If we neglect for a moment the part associated with the confidence parameter β, then, the result simply says that, by extracting a number N of operating conditions as given by (5) and running the corresponding N experiments to evaluate the constraints appearing in (4), the solution (γ_N^*, c_N^*) to (4) violates the constraints corresponding to other, unexperienced, operating conditions with a probability that does not exceed a *user-chosen* level ϵ. This means that the so-determined c_N^* provides an upper bound for the cost $J_{\gamma_N^\star}(\omega)$ valid for every operating condition $\omega \in \Omega$ with the exclusion of at most an ϵ-probability set.

As for the probability $1-\beta$, one should note that (γ_N^*, c_N^*) is a random quantity because it depends on the randomly extracted operating conditions $\omega^{(1)}$, $\omega^{(2)}$, ..., $\omega^{(N)}$. It may happen indeed that these conditions are not representative enough (one could even extract N times the same operating condition!). In this case no generalization is expected, and the fraction of operating conditions violated by (γ_N^*, c_N^*) will be larger than ϵ. Parameter β controls the probability of extracting unrepresentative operating conditions, and the final result that (γ_N^*, c_N^*) violates at most an ϵ-fraction of operating conditions holds with probability $1-\beta$. One important practical fact is that, due to the structure of the equation in (5), β can be set to be so small (say $\beta = 10^{-6}$) that it is virtually zero for any practical purpose, and this does not lead to a significant increase in the value of N (see also the numerical example in Section 4).

For the reader's convenience, the discussion in this section is summarized in a recipe for a practical implementation of the overall adaptive design scheme.

PRACTICAL RESULT: *Select a violation parameter $\epsilon \in (0,1)$, let $\beta = 10^{-6}$, and compute the least integer N satisfying (5). Run N random experiments and compute the corresponding N constraints for problem (4). Then, the solution $\gamma_N^\star$ of (4) achieves performance $c_N^\star$ on all operating conditions but an ϵ fraction of them, and, moreover, $c_N^\star$ is 'better than the best', in the sense that $c_N^\star \leq c^\star$.*

Before closing the section, the following final remark is worth making in the light of equation (5):

> *The number of experiments N that are needed to adapt the device does not depend on the system complexity; it instead only depends on the complexity of the device $\mathcal{D}_\gamma$ through the size k of its parametrization γ.*

Thus reality can be any complex and still we can evaluate the experimental effort by only looking at the device being designed.

4 A numerical example

We consider the problem of inverting the nonlinear characteristic between input u and output $y(u,d)$ of a system affected by an additive output disturbance d (Figure 4), over the range of values $U = [0,1]$ for u (input-output equalization).

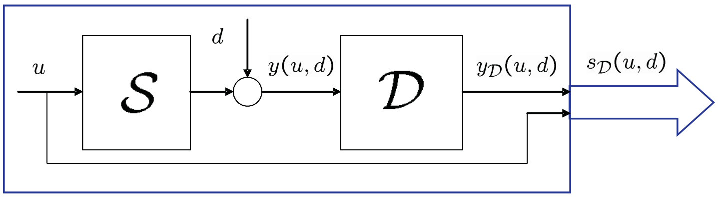

Fig. 4. Inverting a nonlinear characteristic through a device

The device is fed by $y(u,d)$ and produces output $y_{\mathcal{D}_\gamma}(u,d) = \gamma_1 y(u,d)^2 + \gamma_2 y(u,d) + \gamma_3$. The performance of the device with parameter $\gamma = (\gamma_1, \gamma_2, \gamma_3) \in \Re^3$ is given by $\max_{u,d \in U \times D} J_\gamma(u,d)$, where $J_\gamma(u,d) = |y_{\mathcal{D}_\gamma}(u,d) - u|$ and D is the (unknown) range of values for d. In words, this performance expresses the largest deviation off the perfect equalization line $y_{\mathcal{D}} = u$.

We chose $\epsilon = 0.1$, $\beta = 10^{-6}$, and according to (5) N was 205.

The scenario program (4) is in this case

$$\min_{\gamma, c \in \Re^4} c \quad \text{subject to:} \tag{6}$$

$$|\gamma_1 y(u^{(i)}, d^{(i)})^2 + \gamma_2 y(u^{(i)}, d^{(i)}) + \gamma_3 - u^{(i)}| \le c, \; i = 1, 2, \ldots, 205,$$

where $u^{(1)}$, $u^{(2)}$, ..., $u^{(205)}$ are random values for u independently extracted from U according to the uniform distribution P_u over $[0,1]$, and $d^{(1)}$, $d^{(2)}$, ..., $d^{(205)}$ are random values for d independently created by the environment during experimentation according to some (unknown) stationary distribution

P_d.

The 205 constraints in (6) can be evaluated by running 205 experiments on the system where the output samples $y^{(i)} = y(u^{(i)}, d^{(i)})$, $i = 1, 2, \ldots, 205$, are collected together with $u^{(i)}$, $i = 1, 2, \ldots, 205$. Figure 5 shows the outcomes of the experiments. Note that the collected output data present some dispersion due to the presence of the additive disturbance d.

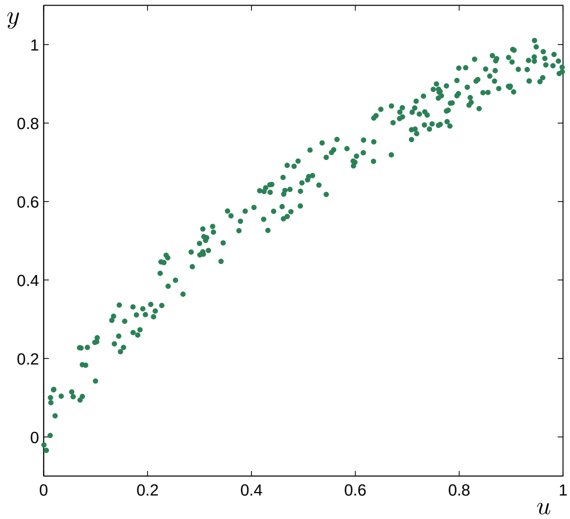

Fig. 5. Outcome of the experiments: samples of input u and output $y(u, d)$

By solving (6) we obtained $\gamma^\star_{205} = (0.424,\, 0.650,\, -0.081)$ and $c^\star_{205} = 0.108$.

$c^\star_{205}$ is the maximum equalization error for the extracted scenarios. In Figure 6, we plot the input and equalized output pairs $(u^{(i)}, y_{\mathcal{D}_{\gamma^\star_{205}}}(u^{(i)}, d^{(i)}))$, $i = 1, 2, \ldots, 205$, and the region $u \pm c^\star_{205} := \{(u, y) : u - c^\star_{205} \leq y \leq u + c^\star_{205},\ u \in U\}$. $u \pm c^\star_{205}$ is the strip of minimum width centered around the perfect equalization line $y_{\mathcal{D}} = u$ that contains all the 205 input and equalized output pairs.

In the light of the practical result at the end of the previous section, device $\gamma^\star_{205}$ carries a guarantee that the equalized output $y_{\mathcal{D}_{\gamma^\star_{205}}}(u, d)$ differs from u of at most $c^\star_{205} = 0.108$ for all u's and d's except for a subset of probability $P = P_u \times P_d$ smaller than or equal to 0.1; moreover, the region of equalization $u \pm c^\star_{205}$ is contained within $u \pm c^\star$. This result holds irrespectively of D and P_d, which are unknown to the designer of the device.

The actual nonlinear characteristic and disturbance d used to generate the data in Figure 5 are shown in Figure 7 together with the designed device

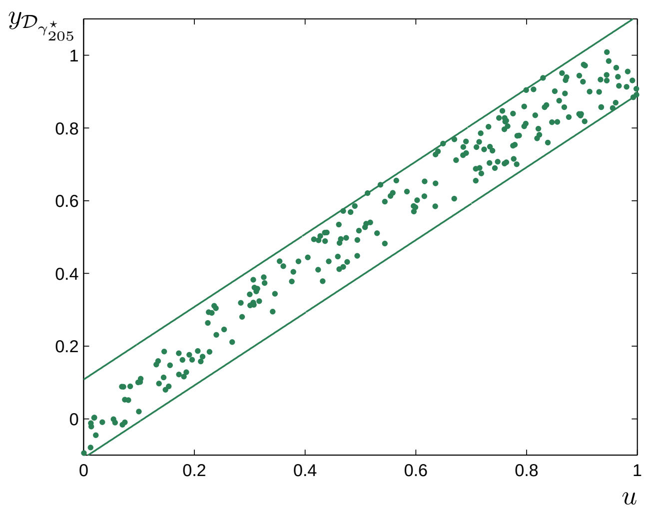

Fig. 6. Input u and equalized output $y_{\mathcal{D}_{\gamma^\star_{205}}}(u,d)$ for the extracted scenarios, and the region of equalization $u \pm c^\star_{205}$

with parameter $\gamma^\star_{205}$. In this example, the parameter of the device could have been designed so as to exactly invert the nonlinear characteristic. However, the obtained $\gamma^\star_{205}$ is different from such a choice, because the device aims at inverting the nonlinear characteristic between u and y while also reducing the effect of d on the reconstructed value for the input u.

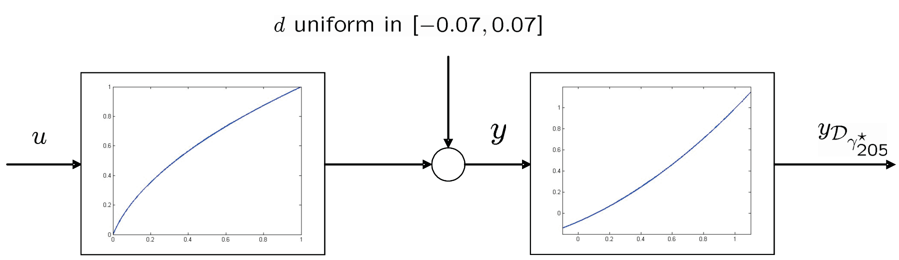

Fig. 7. Actual nonlinear characteristic and disturbance characteristics, along with the designed device

5 Conclusions

The main goal of this contribution is that of attracting the reader's attention to the fundamental issue of evaluating the experimental effort needed to perform adaptive design, and some answers have been provided in a specific worst-case context.

Many are the aspects that our discussion has left unsolved, and open to further investigation:

- it is not always the case that one experiment provides all the information needed to evaluate a constraint. In the disturbance compensator example, for instance, if either the system or the device are not linear it is not possible to swap their order, and constraint evaluation calls for many experiments with virtually all possible devices in place. More generally, more experiments are needed when the input to the system depends on the device being designed.
- a perspective different from the worst-case approach can be used for adaptive design. For example, device quality could be assessed by its average performance, [10, 11, 12], rather than its worst-case performance over the set of operating conditions of interest.

Addressing these problems is a difficult task that requires much additional effort.

References

1. Sastry S., Bodson M. (1994) Adaptive Control: Stability, Convergence, and Robustness. Prentice-Hall.
2. Astrom K.J., Wittenmark B. (1994) Adaptive Control. Addison-Wesley.
3. Bittanti S., Picci G. eds. (1996) Identification, Adaptation, Learning. The science of learning models from data. Springer-Verlag, Berlin, Computer and Systems Science Series, Vol. 153.
4. Landau I.D., Lozano R., M'Saad M. (1998) Adaptive Control. Springer-Verlag.
5. Haykin S. (2002) Adaptive Filter Theory. Prentice Hall.
6. Polderman J.W., Willems J.C. (1998) Introduction to Mathematical Systems Theory: A Behavioral Approach. Springer Verlag, New York.
7. Calafiore G., Campi M.C. (2005) Uncertain convex programs: randomized solutions and confidence levels. Math. Program., Ser. A 102: 25–46.
8. Calafiore G., Campi M.C. (2006) The scenario approach to robust control design. IEEE Trans. on Automatic Control 51(5):742–753.
9. Campi M.C., Garatti S. (2007) The exact feasibility of randomized solutions of convex programs. Internal report, University of Brescia, Italy.
10. Vapnik V.N. (1998) Statistical Learning Theory. John Wiley & Sons.
11. Vidyasagar M. (2001) Randomized algorithms for robust controller synthesis using statistical learning theory. Automatica 37(10):1515-1528.
12. Campi M.C., Prandini M. (2003) Randomized algorithms for the synthesis of cautious adaptive controllers. Systems & Control Letters 49(1):21-36.