Proceedings of the European Control Conference 2007
Kos, Greece, July 2-5, 2007

How much do we need to experiment in order to adapt?

Marco C. Campi

Adaptive control is about constructing controllers formed by two parts, where the first part (estimation unit) is aimed at estimating relevant plant characteristics, and the second part (control unit) exercises the control action by exploiting the outcomes of the estimation part. Adaptation, i.e. control unit tuning, can take place continuously during the normal operation of the system (on-line adaptation), or it can be performed every now and then. The limit case when information is first accrued and used to tune the control unit, which is otherwise kept fixed afterward, is called off-line adaptation.

Concentrating on off-line adaptive schemes, suppose that we can perform experiments on a system and, based on the collected data, we are asked to tune a control unit. One fundamental question to ask is:

> How extensively we need to experiment to be able to come up with a controller of guaranteed properties?

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

A Min-Max perspective. Let $y = P(u, d)$ be a plant, where u is input, y is output, and d is disturbance, and let $u = C_\theta(y, r)$ be a class of controllers parameterized in θ, where r is reference. Also, let $\mathcal{J}(P(u, d), C_\theta(y, r))$ be a control cost. For given P and C_θ, $\mathcal{J}$ depends on d and r, and one can set out to minimizing the worst-case performance of the control system over all possible references and disturbances:

$$\min_\theta \max_{r,d} \mathcal{J}(P(u, d), C_\theta(y, r)).$$

Supposing that the plant P and the disturbance characteristics are not known, how extensively do we need to experiment on the plant in order to come up with a control unit minimizing this $\min\max$ cost? This problem appears to be unanswered at the present stage of knowledge.

A glance at system identification. Referring to a system identification (instead of control) set-up helps shed new light on the above problem.

Consider again $y = P(u, d)$ and let $\hat{y} = \hat{P}_\theta(u)$ be a class of linearly-parameterized models where $\hat{y}$ is the model output which should resembles the system output y. Let the identification cost be: $\mathcal{J}(P(u, d), \hat{P}_\theta(u)) = \sum_t (y_t - \hat{y}_t)^2$, and suppose we want to find θ that attains

$$\min_\theta \max_{u,d} \sum_t (y_t - \hat{y}_t)^2.$$

The interpretation is that the best model is the one that resembles at best the system behavior over all possible operating circumstances. Finding the minimizing θ is a formidable task that requires in principle to experiment on the system in all possible operating conditions. Surprisingly, however, due to recent achievements, a solution close to optimal can be found at a reasonable experimental effort.

An experiment on the system simply consists in randomly selecting an input u, injecting it into the system, and collecting the corresponding output y which will also be affected by the disturbance d. The distribution in the selection of u reflects the intended use of the system. Repeating this experiment, say, N times results in N signals $y^{(1)}, \ldots, y^{(N)}$. One seemingly naive approach for finding θ then consists in substituting the $\max$ in our identification cost with a $\max$ only done over the seen N scenarios:

`scenario optimization:`

$$\min_\theta \max \left\{ \sum_t (y_t^{(1)} - \hat{y}_t^{(1)})^2, \ldots, \sum_t (y_t^{(N)} - \hat{y}_t^{(N)})^2 \right\},$$

where $\hat{y}^{(i)} = \hat{P}_\theta(u^{(i)})$. This minimization can be carried out since it is totally data-based. It is a fact that the resulting θ is nearly optimal for the initial $\min\max$ problem, with no assumptions on P and on the range of variability for d. Precisely, if we assume a stationary set-up through experiments (i.e. plant P is invariant and disturbances d have an invariant probability to happen), the found θ is guaranteed to minimize $\mathcal{J}$ over all possible d and u except at most an ϵ-fraction of them. This is to say that, no matter what P is, in a new unseen situation we have that $\sum_t (y_t - \hat{y}_t)^2 \leq \mathcal{J}_{opt}$ with probability $1 - \epsilon$, where $\mathcal{J}_{opt}$ is the optimal value for the scenario optimization.

Interestingly, ϵ is totally under the user's control. Indeed, one can apply a general formula derived in robust convex optimization, see [1], [2], to evaluate the N that has to be used in order to attain an ϵ-guarantee. This formula scales essentially as $N \sim \frac{size(\theta)}{\epsilon}$ and does not depend on P.

The fact that N does not depend on P, while it does depend on $\hat{P}_\theta$ through $size(\theta)$, has a significant interpretation:

Supported by MIUR under the project "Identification and adaptive control of industrial systems"

Marco C. Campi is with Dipartimento di Elettronica, Università di Brescia, Via Branze 38 – 25123 Brescia, Italy `marco.campi@ing.unibs.it`

ISBN: 978-960-89028-5-5

> It is not important how complex what we try to describe is, what matters is how complex a description of it we want to make.

Back to control. Going back to adaptive control, now we ask: is it possible to establish results that parallel those just mentioned for identification? Indeed, in certain contexts the extension is even straightforward.

Example: noise compensation

Consider the noise compensation scheme in Figure 1. There, C_θ is a linear and linearly-parameterized compensator and P is a linear unknown plant. The objective is to solve the noise-rejection problem

$$\min_\theta \max_d \sum_t y_t^2.$$

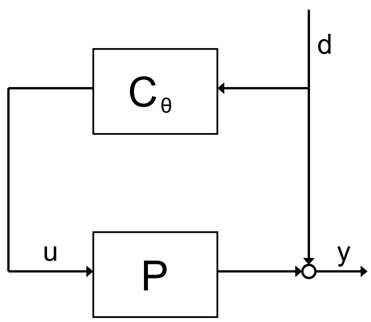

Fig. 1. A compensation scheme

This problem is not much different from the system identification problem. Perform N experiments, where $d^{(i)}$ is measured, and it is directly injected into P; name $y^{(i)} = Pd^{(i)} + d^{(i)}$ the measured output. The output that would have been obtained with the compensator C_θ in place is $PC_\theta d^{(i)} + d^{(i)} = C_\theta P d^{(i)} + d^{(i)} = C_\theta(y^{(i)} - d^{(i)}) + d^{(i)}$ and it can be computed from the measured signals $y^{(i)}$ and $d^{(i)}$. Thus, one can set out to solving the scenario counterpart of the noise-rejection problem:

$$\min_\theta \max \left\{ \sum_t (C_\theta(y_t^{(1)} - d_t^{(1)}) + d_t^{(1)})^2, \ldots, \right.$$
$$\left. \sum_t (C_\theta(y_t^{(N)} - d_t^{(N)}) + d_t^{(N)})^2 \right\}.$$

Similarly to identification, what determines the number N of experiments needed in order that the obtained θ is a ϵ-guaranteed solution for the initial noise-rejection problem is the compensator complexity, not the complexity of reality. $*$

The previous example is very specific in many respects: plant and controller are linear and there is no feedback. Investigating more general situations – in the same vain of analysis as done here or, possibly, along other approaches – to assess the experimental effort required to learn the environment to perform adaptive design represents a fundamental theoretical challenge.

References

[1] G. Calafiore and M.C. Campi, The scenario approach to robust control design, *IEEE Trans. on Automatic Control*, vol. 51, 2006, pp 742–753.

[2] G. Calafiore and M.C. Campi, Uncertain convex programs: randomized solutions and confidence levels, *Mathematical Programming*, vol. 102, 2005, pp 25–46.