The local approach to change detection, diagnosis, and model validation: application to vibration mechanics*

Albert Benveniste, Michèle Basseville, Maurice Goursat, and Laurent Mével

To Lennart, who donated, to the system identification galaxy, one new algorithm for each drank bottle of coke.

Abstract: In this paper we survey our results on model validation, change detection, and diagnosis, based on the *local approach.* We see this work as a counterpart of the central contribution by Lennart Ljung to a systematic approach to system identification. We reuse concepts brought by Lennart such as: model set, true system model, identification method, and we enrich this set of concepts by considering also the nominal model.

The local approach consists in assuming that nominal and true models differ by a factor of order $1/\sqrt{N}$, where N is the sample length. This allows deriving Gaussian approximations for the (generalized) likelihood ratio statistics relating these two models. This allows designing systematically criteria and test statistics for model validation, change detection, and diagnosis.

We report our experience in using these techniques in the area of vibrations monitoring in mechanical engineering, where they have proven very useful and effective.

1 Introduction: why Lennart Ljung made this work possible

As very well explained by Michel Gevers Gevers (2005), the essential and deepest contribution of Lennart Ljung to system identification was the clarification of the following key concepts:

- the data and the system producing the data;
- the model set and the best fit within the model set;

*This work has been supported by Eureka projects 1562 Sinopsys, 2419 Flite, and 3341 Flite2, and by project ACI S&I CONSTRUCTIF, funded by the French minitère de la recherche.

- the algorithm to estimate the best fit and its bias and variance error.

All this occurred in the late seventies and late 70's Ljung (1977b,a, 1978); Ljung and Soderström (1983), and contributed dramatically to the clarification of the area. It considerably cleaned the jungle of system identification algorithms, by classifying the variants with respect to each of the above concepts.

In the late seventies, our group started working on change detection. At that time the area was dominated by results from statisticians Brown et al. (1975); Hinkley (1970); Lehmann (1968); Shiryaev (1961); Wald (1947); Willsky and Jones (1967). While these were very elegant and technically deep results, they were essentially focusing on likelihood based methods.

Sometimes problems are raised that cause decisive breaks in the activity of researchers. Our group experienced this when Bruno Barnouin, from CNEXO (now IFREMER, the French institute for research on sea), asked us in 1981 to consider the problem of detecting fatigues in offshore structures before the damage could actually occur. One difficulty was that measured data were clearly nonstationary, much more so because of the turbulent nature of the excitation by the swell and cavitation effects that follow, than because of the damage itself. We were able to show that Instrumental Variable methods for output only eigenstructure identification were robust against such input nonstationarities Benveniste and Fuchs (1985).

This opened the possibility to perform fatigue detection through the comparison of results from identifications, performed on the one hand while the structure was safe, and on the other hand at the current instant. This was not satisfactory to us, however, since identification of structures subject to vibrations was (and still is) difficult and not fully automatic. We thought that deciding upon changes should, by principle, be simpler than, first, performing identification, and, second, comparing the models.

Two sources of inspiration helped us finding what we wanted. First, Igor Nikiforov Nikiforov (1983) made us aware that Le Cam, a statistician, discovered in the late sixties that it was indeed natural to consider testing problems in which the alternatives were closer along with the size of the data set at hand: the larger the data set, the closer the alternatives for discrimination can be Le Cam (1960); Davies (1973); Le Cam (1986); Roussas (1972). Taking a gap of order $1/\sqrt{N}$ (where N is the size of the data set) gave a likelihood ratio that was asymptotically Gaussian. This was called the *local hypothesis* and Le Cam's approach and its elegant consequences were called the *local approach.*

The original Le Cam local approach was ideologically "likelihoodist", however. We knew we wanted to work with Instrumental Variables, for their robustness properties. Ljung and the swedish school had paved the way we should go: clarify the concepts, and then go for a possible generalization of the likelihood techniques. Once we realized this, finding the right concepts was not too dif-

ficult. The ODE concept "$\dot{\theta} = h(\theta)$" had demonstrated the importance of the mean gradient $h(\theta)$; the asymptotic behavior of the ODE associated with recursive stochastic algorithms could tell us a lot about the asymptotic behavior of the algorithm itself. For our change detection problem, the right concepts became clear:

- the data and the (unknown) system $\theta_\star$ producing the data;
- the nominal model θ_0 to confront to the data;
- the statistics to assess the above confrontation, by using the nominal model and the data at hand; corresponding to Ljung's $h(\theta)$ function was a function $h(\theta_0, \theta_\star)$ assessing "in the mean" the gap between θ_0 and $\theta_\star$.

This was developed in Benveniste et al. (1987, 1990); Basseville and Nikiforov (1993); Zhang et al. (1994). Still, this approach remained essentially confidential, much more than Lung's results for identification. It is only in the area of modal identification and health monitoring of structures in mechanical engineering that this approach found an interesting field of application. Reasons for this possibly are: data sets are large, models are huge (several dozens of poles, corresponding to reduced order models for FE models with thousands of degrees of freedom), and damages results in slight changes in the model characteristics (a fraction of a percent to a few percent for the frequencies).

This is why we have chosen to collect in this paper in honor of Lennart Ljung our experience in modal monitoring with local approaches.

2 Running application: modal identification and health monitoring of mechanical structures

In this section we discuss our application area, namely modal identification and health monitoring of mechanical structures.

2.1 Context and benchmark example

The design and maintenance of mechanical structures subject to noise and vibrations is an important topic in mechanical engineering. It is an important component of comfort (cars and buildings) and contributes significantly to the safety related aspects of design and maintenance (aircrafts, aerospace vehicles and payloads, civil structures). Requirements from these application areas are numerous and demanding:

- Detailed physical models derived from first principles are developed as part of system design (although detailed modeling of the links with the ground

is not feasible, for civil engineering structures). These models involve the dynamics of vibrations, sometimes complemented by other physical aspects (fluid-structure interaction, aerodynamics, thermodynamics).

- Laboratory and in-operation tests are performed on the prototype structure, in order to get so-called *modal models*, *i.e.,* to extract the modes and damping factors (these correspond to system poles), the mode shapes (corresponding eigenvectors), and loads. These results are used for *updating* the design model for a better fit to data, and sometimes for certification purposes (*e.g.,* in flight domain opening for new aircrafts).

- The *monitoring* of structures is an important activity for the system maintenance and health monitoring. This is particularly important for civil structures. Damaged structures would typically exhibit changes in their stiffness due to the occurrence of micro-cracks or faults, or modifications of the links to the ground. A key difficulty is that such system characteristics are also sensitive to environmental conditions, such as temperature effects (for civil structures), or external loads (for aircrafts). In fact these environmental effects even dominate the effect of damage. This is why, for very critical structures such as aircrafts, detailed active inspection of the structures is performed as part of the maintenance. Still, in many cases (*e.g.,* civil structures), health monitoring based on modal techniques is the preferred, non invasive, mode of monitoring. Of course, *the localization of a damage must be expressed in terms of the physical model, not in terms of the modal model used in system identification.*

Consequently, the following elements are encountered and must be jointly dealt with when addressing these applications: *design models* from the system physics, *modal models* used in structural identification, and, of course, *data* from sensors. Corresponding characteristics are given now:

- Design models are Finite Element models, with tens or hundreds of thousands elements. These models are linear if only small vibrations are considered; still, these models are of huge order. In addition, nonlinearities enter as soon as large vibrations or other physical effects (aerodynamics, thermodynamics...) are considered.

- Sensors can range from a handful of accelerometers or strain gauges, to hundreds of them, or even more if laser based measurement technology is used.

- Consequently, modal models used for structural identification are of very high order, and are always the result of drastic model reduction.

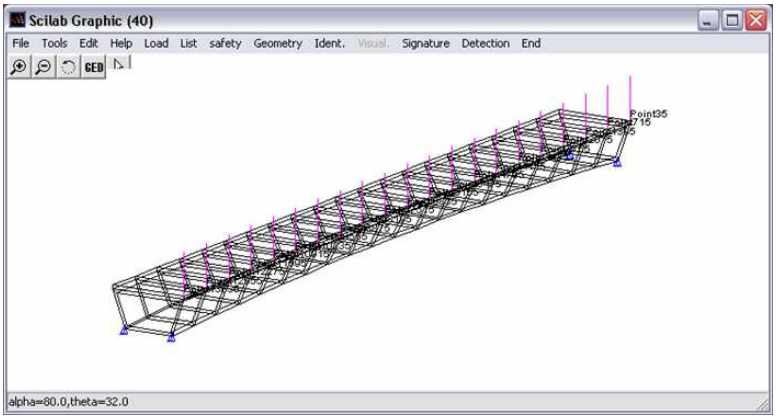

Figure 1: *The simulation example: a realistic bridge (courtesy of Etienne Balmes).*

Fig. 1 shows a simulated bridge structure (courtesy of Etienne Balmes), which will support our various use cases. The structure has 13668 nodes, 9642 finite elements, and 40976 degrees of freedom. Its first five modes are listed in Table 1.

Mode	f_1	f_2	f_3	f_4	f_5
Safe	1.28	2.63	2.7	6.05	6.34
Damaged	1.29	2.61	2.69	6.05	6.33

Table 1: *First five modes of the structure of Fig. 1 (in Hz).*

We now briefly indicate which parameters are used for identification. For time invariant structures subject to small vibrations, the classical discrete time linear model can be used, namely:

$$\left\{ \begin{array}{lcl} X_{k+1} & = & AX_k + V_k \\ Y_k & = & CX_k + W_k \end{array} \right. \tag{1}$$

The *modes* of the structure are nothing but the poles of the system, *i.e.*, the eigenvalues of A. Mode information decomposes into a *frequency* and a *damping.* For λ a mode, its associated *mode shape* is $\varphi_\lambda =_{\rm def} C\Phi_\lambda$, where Φ_λ is the eigenvector of A corresponding to λ. Intuitively, the mode shape φ_λ indicates, in the complex domain, how the structure behaves, for its motion at frequency ω_λ. The collection of modes and mode shapes constitutes the eigenstructure of the system, they are collected in the parameter vector θ. Eigenstructure θ is a canonical

parameterization of the pole part of system (1). In mechanical engineering jargon, this is referred to as the *modal model* of the structure. All use cases we shall discuss handle the modal model with or without joint use of the finite element model.

2.2 Some new difficulties as compared to system identification

We argued in the introduction that detection is easier than identification, so that one should not perform re-identification in order to detect and diagnose changes. In this section, we shall see that, on the other hand, change detection raises new problems, not encountered in system identification. We discuss these in the context of modal monitoring.

2.2.1 Model matching when models are close

We first discuss a problem that is typically encountered in modal monitoring. Assume that the designer has a nominal model for the above system, together with a measured data set. The nominal model can be, for instance, obtained from a Finite Element design model. The designer wishes to focus on the frequency range 1–8Hz, where this nominal model possesses five poles denoted by $f_{0,1}, \ldots, f_{0,5}$, of which only $f_{0,3}, f_{0,4}, f_{0,5}$ are identifiable with the selected sensor positions. The designer wishes to know if these three poles match her data set well.

One possible way or performing this would be as follows: the designer would perform modal identification on the measured data set, and then compare the obtained modes with the nominal ones. This is illustrated on Fig. 2. On the top diagram of this figure, we zoom on a restricted bandwidth of a plot of poles identified for this data set, for different model orders. This is called a *stabilization diagram* by mechanical engineers. Stabilization diagrams are used to (visually or automatically) select relevant poles in a very high order system, *i.e.*, poles that are stable for different model orders and have mechanical meaning. The three lines of poles correspond to f_3, f_4, and f_5, ordered from bottom to top.

Now, comparing the nominal poles with the stabilization diagram of Fig. 2 raises the difficulty of *model matching*: how to pair the nominal poles with the identified ones, especially when poles are close? In contrast, Fig. 2, bottom diagram, shows that having a good model/data gap allows to assess how good the nominal model is, without being faced with the model matching problem.

2.2.2 Nuisance parameters

Fig. 3 shows another, more important and more difficult, problem. Vibrating structures are subject to damages and temperature effects. Both factors affect

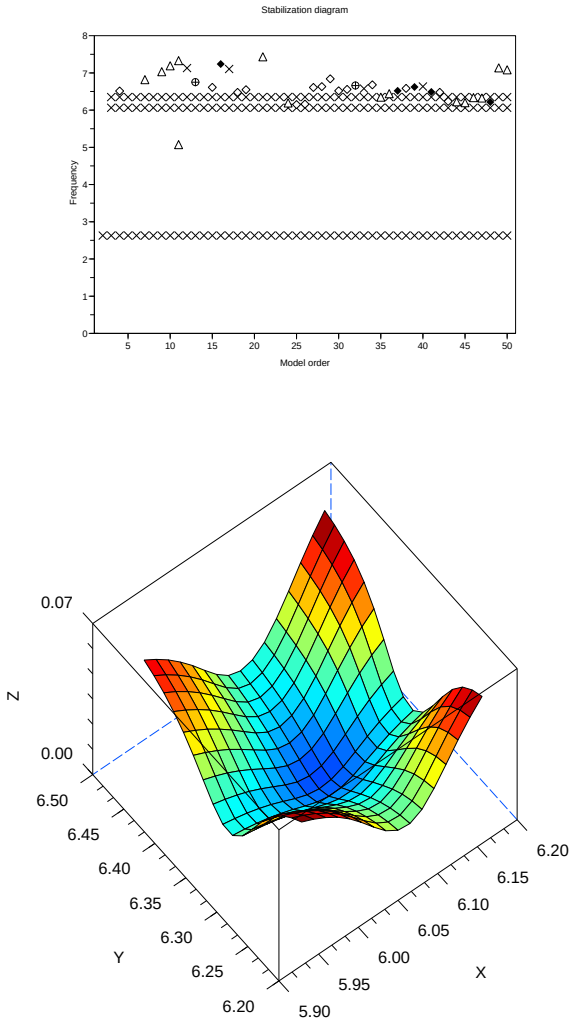

Figure 2: *Simulation example of Fig. 1. Top: stabilization diagram showing three frequencies for different model orders. The three frequencies are* f_3, f_4, f_5*, ordered from bottom to top. Bottom: model/data gap when the two frequencies* f_4 *(on X-axis), and* f_5 *(on Y-axis), is shifted below and above its nominal value by* 2.5%*.*

their dynamics, but only damage is for monitoring by the engineer. In Fig. 3, we

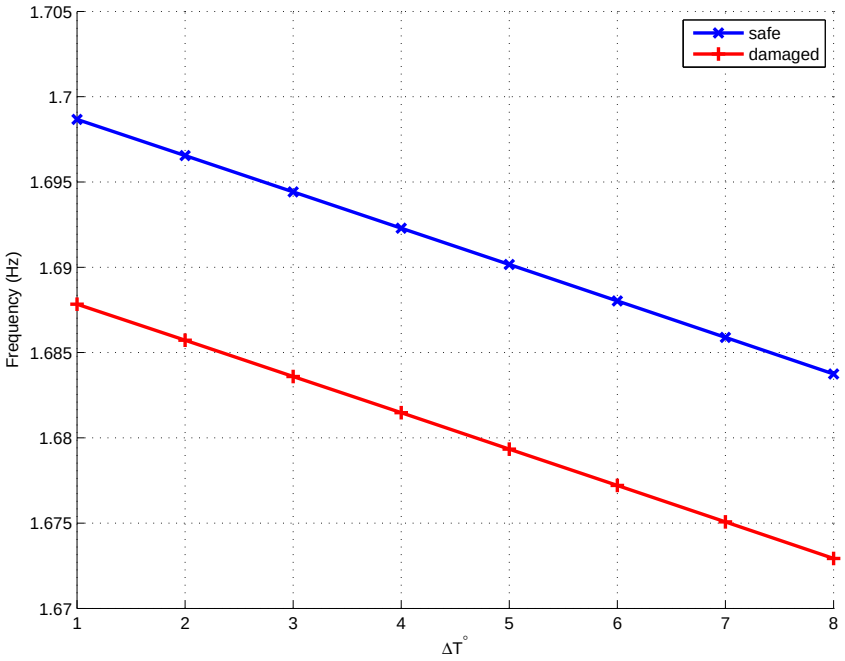

Figure 3: *Simulation example of Fig. 1. We plot how a given system frequency is shifted on a structure subject to vibrations when the temperature increases from* 0 *to* 8 *degrees Celsius, for a safe (in blue) and damaged (in red) structure, respectively.*

plot how one given pole is shifted when the temperature increases, from 0 to 8 degrees Celsius, for a safe (in blue) and damaged (in red) structure, respectively. As the reader can see, for an increase of 6 degrees and more, the temperature effect is larger than that of the damage at 0 degree Celsius.

This is referred to as *nuisance* parameters, *i.e.*, parameters that cause changes in the observed dynamics but are irrelevant for the problem in consideration.

3 The local approach for the independent, identically distributed case, with likelihood based methods

3.1 Problem setting

In this section, we assume a sequence Y_k of independent, identically distributed (i.i.d.) random variables whose distribution $p_\theta(Y_k)$ depends on some parameter vector θ. Symbol $\mathbf{E}_\theta$ denotes the expectation under distribution p_θ, and "ln" denotes the logarithm. For a given N-size sample $\mathcal{Y}_1^N =_{\rm def} Y_1, \ldots, Y_N$, two values of interest for parameter vector θ are considered:

- The *true* value $\theta_\star$ is such that Y_k was actually drawn under $p_{\theta_\star}$; the true value is generally not known.
- A *nominal* value θ_0 that is given, regardless of the considered sample. Typically, θ_0 may be chosen by the user; or it may be given by some computed design model; or it may be the true or estimated value for another data sample.

The aim is to confront θ_0 to $\theta_\star$, or, equivalently, to confront θ_0 to the observed data sample $\mathcal{Y}_1^N =_{\rm def} Y_1, \ldots, Y_N$.

In the sixties, Le Cam made the following observation: the longer the available data set is, the closer models we wish to discriminate can be. Thus it makes sense assuming a deviation of order $1/\sqrt{N}$ between nominal and true models, where N is the sample length:

$$\theta_{_0} - \theta_\star \;=\; \frac{1}{\sqrt{N}}\,\delta_{,}. \tag{2}$$

In (2), δ is a given fixed vector "gap" not depending on sample size. Normalization (2) is known as the *statistical local approach*. Under (2), a number of simplifications occur that can be exploited for our present purpose.

3.2 A simple example where calculations can be made explicit

As a simple example, consider the case of scalar i.i.d. Gaussian random variables $y_k \sim \mathcal{N}(\theta, \sigma^2)$, having fixed and known variance σ^2 but unknown translation parameter θ. Assume that $\theta_\star = 0$ and compute the likelihood and negative loglikelihood, respectively:

$$p_\theta(y) = \frac{1}{\sigma\sqrt{2\pi}}\, e^{-\frac{(y-\theta)^2}{2\sigma^2}} \;,\; -\log p_\theta(y) = \log(\sigma\sqrt{2\pi}) + \frac{(y-\theta)^2}{2\sigma^2}.$$

Next, consider the following quantity:

$$\zeta_N(\theta) \ =_{\rm def} \ \frac{1}{\sqrt{N}} \sum_1^N \frac{\partial}{\partial\theta} \log p_\theta(y_k) \ = \ \frac{1}{\sqrt{N}} \sum_1^N \frac{y_k - \theta}{\sigma^2}$$

Using the local approach (2), renormalize $\theta_{_0} - \theta_\star = \frac{1}{\sqrt{N}}\,\delta$; this yields

$$\zeta_N(\theta_{_0}) \ = \ -\frac{\delta}{\sigma^2} + \left(\frac{1}{\sqrt{N}} \sum_1^N \frac{y_k}{\sigma^2} \right)$$

whence, for N large, by Central Limit Theorem,

$$\boxed{\zeta_N(\theta_{_0}) \ \sim \ -\frac{\delta}{\sigma^2} + \mathcal{N}\left(0, \frac{1}{\sigma^2}\right)} \tag{3}$$

Data-to-model gap $\zeta_N(\theta)$ is proportional to signal/noise ratio and subject to a $1/\sigma^2$ random perturbation. In the next section we generalize this to any vector i.i.d. random sequence Y_k.

3.3 Basic facts supporting the local approach

The *log-likelihood* of $\mathcal{Y}_1^N \ =_{\rm def} \ (Y_1, \ldots, Y_N)$ is defined as:

$$\ln p_\theta(\mathcal{Y}_1^N) \ =_{\rm def} \ \sum_{k=1}^N \ l_\theta(Y_k)\,, \ \text{where } l_\theta(y) \ =_{\rm def} \ \ln p_\theta(y)\,.$$

For any given θ_0, the *efficient score* $\zeta_N(\theta_0)$ is the value, at θ_0, of the gradient, w.r.t. θ, of the log-likelihood normalized by $\sqrt{N}$:

$$\zeta_N(\theta_0) \ =_{\rm def} \ \frac{1}{\sqrt{N}} \ \frac{\partial}{\partial\theta} \ln p_\theta(\mathcal{Y}_1^N)\Big|_{\theta=\theta_0} \ = \ \frac{1}{\sqrt{N}} \ \sum_{k=1}^N \frac{\partial}{\partial\theta} l_\theta(Y_k)\Big|_{\theta=\theta_0} \tag{4}$$

Note that the efficient score only depends on the data and the nominal model. It is therefore available to the designer.

Next, the following function of the pair consisting of the nominal and true models will be important:

$$h(\theta_0, \theta_\star) \ =_{\rm def} \ \mathbf{E}_{\theta_\star} \left(\frac{\partial}{\partial\theta} l_\theta(Y_k)\Big|_{\theta=\theta_0} \right)$$

Observe that the function inside the parentheses is parameterized by the nominal model θ_0, whereas expectation holds under the true model $\theta_\star$ (which is the distribution of the data at hand).

Consequences of assuming $\theta_0 = \theta_\star$ *i.e.,* the nominal model equals the true one. For this paragraph, we assume $\theta_0 = \theta_\star$ and discuss the consequences. Since $\mathbf{E}_\theta(Y_k) = \int p_\theta(y)dy = 1$ holds for every θ, condition $\theta_0 = \theta_\star$ implies:

$$h(\theta_0, \theta_\star) = \int \Big(\frac{\partial}{\partial\theta} l_\theta(y)\Big|_{\theta=\theta_0}\Big) p_{\theta_\star}(y)\, dy = \Big(\frac{\partial}{\partial\theta}\int p_\theta(y)dy\Big)\Big|_{\theta=\theta_\star} = 0. \quad (5)$$

Formula (5) has the following consequences. First, the efficient score is unbiased:

$$\text{for every } N,\ \theta_0 = \theta_\star \text{ implies } \mathbf{E}_{\theta_\star}(\zeta_N(\theta_0)) = 0\,. \quad (6)$$

Second, by the Law of Large Numbers (LLN) and formula (5), $\theta_0 = \theta_\star$ implies

$$\frac{1}{\sqrt{N}}\zeta_N(\theta_0) \;=\; \frac{1}{N}\sum_{k=1}^{N} \frac{\partial}{\partial\theta} l_\theta(Y_k)\Big|_{\theta=\theta_0} \longrightarrow 0 \text{ when } N \to +\infty. \quad (7)$$

Third, by the Central Limit Theorem, if $\theta_0 = \theta_\star$ holds, then

$$\zeta_N(\theta_0) \;\to\; \mathcal{N}\Big(0\,,\, \mathbf{I}(\theta_\star)\Big) \text{ when } N \to +\infty, \quad (8)$$

where $\mathbf{I}(\theta)$ is the *Fisher information matrix* defined by:

$$\mathbf{I}(\theta) \;=_{\text{def}}\; \mathbf{E}_\theta\; \left[\left(\frac{\partial}{\partial\theta} l_\theta(Y_k)\right)\left(\frac{\partial}{\partial\theta} l_\theta(Y_k)\right)^T\right] = -\,\mathbf{E}_\theta\left[\frac{\partial^2 l_\theta(Y_k)}{\partial\theta^2}\right]. \quad (9)$$

Using Le Cam's *local approach*. Under the local approach (2), taking a first order Taylor expansion of the efficient score around $\theta_\star$ yields:

$$\zeta_N(\theta_0) - \zeta_N(\theta_\star) \;\approx\; \left(\frac{1}{N}\sum_{k=1}^{N} \frac{\partial^2 l_\theta(Y_k)}{\partial\theta^2}\Big|_{\theta=\theta_\star}\right)\, \delta \;\approx\; -\mathbf{I}(\theta_\star)\,\delta, \text{ by (9).} \quad (10)$$

By (8) and (10), the local hypothesis (2) implies:

$$\boxed{\zeta_N(\theta_0) \;\to\; \mathcal{N}\Big(-\mathbf{I}(\theta_\star)\,\delta\,,\, \mathbf{I}(\theta_\star)\Big) \text{ when } N \to +\infty.} \quad (11)$$

In (11), it turns out that, for N large and under the local hypothesis (2), we can equally well evaluate the Fisher matrix at $\theta_\star$ or at θ_0. This flexibility is important in practice.

The efficient score as a data/model gap. We shall use (11) in many ways throughout this paper, by interpreting the efficient score as a data/model gap.

3.4 Using the local approach for model validation or change detection

Model validation can be defined as the following problem:

Problem 1 (model validation). We are given a *nominal model* θ_0 and a N-size data set $\mathcal{Y}_1^N$. Can we accept or should we rather reject the hypothesis that $\mathcal{Y}_1^N$ was drawn under distribution p_{θ_0}?

Under the local approach, this can be answered by using (11) as follows: simply replace Problem 1 with the asymptotically equivalent

Problem 2 (model validation, local equivalent). We are given a *nominal model* θ_0 and the efficient score $\zeta_N(\theta_0)$. Can we accept or should we rather reject the hypothesis that $\delta = 0$ in the normal distribution (11) of $\zeta_N(\theta_0)$?

The interest of Problem 2 is that it is much simpler to solve than Problem 1, since it involves deciding whether or not a Gaussian vector random variable of known covariance matrix has zero mean. The relevant test is based on the following χ^2-statistics:

$$\boxed{\chi(\theta_0) \;=_{\text{def}}\; \zeta_N^T(\theta_0)\,\mathbf{I}^{-1}(\theta_\star)\,\zeta_N(\theta_0),} \tag{12}$$

which must be distributed according to a centered χ^2 with a known number of degrees of freedom if $\delta = 0$ holds—this property can be used to tune the threshold when deciding whether or not $\delta = 0$ can be accepted.

Still, an issue remains. Since $\theta_\star$ in (12) is unknown, so is the matrix $\mathbf{I}(\theta_\star)$. However, by (10) and subsequent remark, this unknown matrix $\mathbf{I}(\theta_\star)$ can be replaced in (12) by its estimator

$$\widehat{\mathbf{I}}_N(\theta_0) \;=_{\text{def}}\; \left(\frac{1}{N} \sum_{k=1}^{N} \left. \frac{\partial^2 l_\theta(Y_k)}{\partial \theta^2} \right|_{\theta=\theta_0} \right)$$

which is based only on the known nominal model and observed data. Our final statistics to answer Problem 2, and hence also (asymptotically) Problem 1, is therefore

$$\chi_N(\theta_0) \;=_{\text{def}}\; \zeta_N^T(\theta_0)\,\widehat{\mathbf{I}}_N^{-1}(\theta_0)\,\zeta_N(\theta_0)\,, \tag{13}$$

which has to be compared against some pre-tuned threshold—this threshold is tuned to achieve a pre-specified level for the test. The χ^2 test statistics (13) takes into account both the size of the deviation of the nominal model from true model, as well as all uncertainties due to measurement noises and other type of randomness.

3.5 Using the local approach for system identification

Here we consider the problem of identifying the true model $\theta_\star$ from an N-size data set. To this end, we consider the *maximum likelihood estimator* $\widehat{\theta}_N$ defined by

$$\boxed{\widehat{\theta}_N \;=_{\rm def}\; \arg_\theta \left[\frac{1}{N}\frac{\partial}{\partial\theta}\ln p_\theta(\mathcal{Y}_1^N) \;=\; 0\right] \;=\; \arg_\theta\left[\zeta_N(\theta) \;=\; 0\right]} \tag{14}$$

By definition of $\widehat{\theta}_N$, $\frac{1}{N}\sum_1^N \frac{\partial}{\partial\theta} l_\theta(Y_k)\big|_{\theta=\widehat{\theta}_N} = 0$. On the other hand, by the law of large numbers, for N large, $\frac{1}{N}\sum_1^N \frac{\partial}{\partial\theta} l_\theta(Y_k)\big|_{\theta=\widehat{\theta}_N} \approx h(\widehat{\theta}_N, \theta_\star)$. [1] Hence

$$0 \approx h(\widehat{\theta}_N, \theta_\star) \text{ for } N \text{ large.} \tag{15}$$

Next, focus on (5). It says that $\theta = \theta_\star$ implies $h(\theta, \theta_\star) = 0$. Assume, conversely, that $h(\theta, \theta_\star) = 0$ implies $\theta = \theta_\star$ (this holds under suitable regularity conditions not detailed here). Now, with this assumption, using (15) and assuming h regular yields that *the maximum likelihood estimator $\widehat{\theta}_N$ converges to the true parameter $\theta_\star$ when $N \to +\infty$* Delyon et al. (1997).

3.6 Relating system identification and model validation

Applying (2) and (10) with $\theta_0 = \widehat{\theta}_N$ and recalling (14), yields

$$\boxed{\sqrt{N}\,\left(\widehat{\theta}_N - \theta_\star\right) \;\approx\; \mathbf{I}^{-1}(\theta_\star)\,\zeta_N(\theta_\star)} \tag{16}$$

Eqn. (16) relates the identification error $\widehat{\theta}_N - \theta_\star$ to the efficient score $\zeta_N(\theta_\star)$ at the true model, under the local approach. Informally, Eqn. (16) relates identification to model validation. On the other hand, the asymptotic distribution of $\zeta_N(\theta_\star)$ is known to us: just take $\delta = 0$ in (11). Combining this observation with (16), we re-derive the well known fact that the maximum likelihood estimator has asymptotic convergence rate

$$\boxed{\sqrt{N}\,\left(\widehat{\theta}_N - \theta_\star\right) \;\approx\; \mathcal{N}\left(0\,,\,\mathbf{I}^{-1}(\theta_\star)\right)} \tag{17}$$

[1] An expert reader may have recognized that we cheated somehow at this point: the LLN holds for a *fixed* value for parameter θ; here we used it for a varying one, since $\widehat{\theta}_N$ depends on the data. This is kind of a "uniform" LLN that holds under suitabe regularity conditions Delyon et al. (1997), not detailed here.

4 The local approach for the general case with arbitrary methods

In Section 3, we have discussed the local approach for i.i.d. sequences based on likelihood techniques, and we have seen that both topics are indeed tightly related. This is nice but still not enough as this assumes that: 1/ likelihood methods are applied, and 2/ the efficient score can be effectively computed. The latter is a demanding assumption, since likelihood functions are not easily computed, except for specific classes of distributions. It is therefore tempting to look for a generalization of the above approach to more general distributions and methods.

In this section, we thus assume a stationary sequence Y_k of possibly dependent random variables whose *joint* distribution $p_\theta(\mathcal{Y}_1^N)$ depends on some parameter vector θ, for each sample size N. Symbol $\mathbf{E}_\theta$ denotes the expectation under this joint distribution. Fortunately, the local theory generalizes to this more general setting, under appropriate smoothness assumptions, see Delyon et al. (1997). We now give a brief presentation of this.

4.1 The pseudo-score

For our general case, we assume that the counterpart of the efficient score is available in the form of the following *pseudo-score*

$$\zeta_N(\theta) \;=_{\rm def}\; \frac{1}{\sqrt{N}} \sum_{k=1}^{N} H(\theta, Y_k)$$

where the function $H(\theta, y)$, is such that—compare with (5):

$$h(\theta_0, \theta_\star) \;=_{\rm def}\; \mathbf{E}_{\theta_\star}\left(H(\theta_0, Y_k)\right) \quad \text{satisfies:} \quad \theta_0 = \theta_\star \;\Leftrightarrow\; h(\theta_0, \theta_\star) = 0 \tag{18}$$

By the LLN for stationary random sequences, condition $\theta_0 = \theta_\star$ implies:

$$\frac{1}{\sqrt{N}}\zeta_N(\theta_0) = \frac{1}{N}\sum_{k=1}^{N} H(\theta_0, Y_k) \quad \longrightarrow \quad 0, \tag{19}$$

and, by the Central Limit Theorem, again for stationary random sequences,[2] if $\theta_0 = \theta_\star$ holds, then

$$\zeta_N(\theta_0) \quad \longrightarrow \quad \mathcal{N}\left(0\,,\, \mathbf{R}(\theta_\star)\right) \tag{20}$$

[2] CLT for stationary but dependent random sequences are not easy to obtain; typically some kind of asymptotic independence is assumed, meaning that observations separated by a long period of time are almost independent. See Benveniste et al. (1990); Delyon et al. (1997) for details.

for some covariance matrix $\mathbf{R}(\theta_\star)$. Closed form formulas exist for $\mathbf{R}(\theta_\star)$ but are not very useful. Empirical estimates for this covariance matrix, based on random trials of $\zeta_N(\theta_0)$, can be used instead—this is referred to as *subsampling* in statistics Davison and Hinkley (1997).

4.2 Extending the local approach

Regarding the generalization of the local approach, the following can be said. Taking a first order Taylor expansion of $\zeta_N(\theta)$ around $\theta_\star$ yields that, under the local hypothesis

$$\theta_0 - \theta_\star \;=\; \frac{1}{\sqrt{N}}\,\delta,$$

we have

$$\zeta_N(\theta_0) - \zeta_N(\theta_\star) \;\approx\; \left(\frac{1}{N}\sum_{k=1}^{N}\left.\frac{\partial H(\theta, Y_k)}{\partial\theta}\right|_{\theta=\theta_\star}\right)\,\delta \;\approx\; -\mathbf{J}(\theta_\star)\,\delta, \qquad (21)$$

where

$$\mathbf{J}(\theta) \;=_{\rm def}\; -\mathbf{E}_\theta\left(\frac{\partial H(\theta, Y_k)}{\partial\theta}\right) \qquad (22)$$

Now, combining (20) and (21) yields (compare with (11) and observe in particular that, in general, $\mathbf{J}(\theta) \neq \mathbf{R}(\theta)$): under local hypothesis $\theta_0 - \theta_\star \;=\; \frac{1}{\sqrt{N}}\,\delta$, we have

$$\boxed{\zeta_N(\theta_0) \;\rightarrow\; \mathcal{N}\Big(-\mathbf{J}(\theta_\star)\,\delta\,,\,\mathbf{R}(\theta_\star)\Big) \text{ when } N \rightarrow +\infty.} \qquad (23)$$

which yields the following generalization for the chi-square test introduced in (12):

$$\boxed{\chi(\theta_0) \;=_{\rm def}\; \zeta_N^T(\theta_0)\,\mathbf{R}^{-1}\mathbf{J}\left(\mathbf{J}^T\mathbf{R}^{-1}\mathbf{J}\right)^{-1}\mathbf{J}^T\mathbf{R}^{-1}\,\zeta_N(\theta_0),} \qquad (24)$$

Regarding identification procedures, so-called M-estimators have been first considered in the sixties by Huber Huber (1981)—compare with (14):

$$\boxed{\widehat{\theta}_N \;=_{\rm def}\; \arg_\theta\left[\,\zeta_N(\theta) \;=\; 0\right].} \qquad (25)$$

Applying (21) with $\theta_0 = \widehat{\theta}_N$, and recalling (25) yields (we assume $\mathbf{J}$ invertible for simplicity):

$$\boxed{\sqrt{N}\,\left(\widehat{\theta}_N - \theta_\star\right) \;\approx\; \mathbf{J}^{-1}(\theta_\star)\,\zeta_N(\theta_\star)} \qquad (26)$$

i.i.d. / likelihood	**general / pseudo-score**
$H(\theta, y) = \frac{\partial}{\partial\theta} \log p_\theta(y)$	select $H(\theta, y)$ such that
$h(\theta_0, \theta_\star) = \mathbf{E}_{\theta_\star} H(\theta_0, Y_k)$	$h(\theta_0, \theta_\star) = \mathbf{E}_{\theta_\star} H(\theta_0, Y_k)$; this implies
$h(\theta_0, \theta_\star) = 0 \Leftrightarrow \theta_0 = \theta_\star$	$h(\theta_0, \theta_\star) = 0 \Leftrightarrow \theta_0 = \theta_\star$
score $\zeta_N(\theta)$	pseudo-score $\zeta_N(\theta)$
$=_{\rm def}\ \frac{1}{\sqrt{N}} \sum_1^N H(\theta, Y_k)$	$=_{\rm def}\ \frac{1}{\sqrt{N}} \sum_1^N H(\theta, Y_k)$
$\theta_0 - \theta_\star = \frac{\delta}{\sqrt{N}}$ implies	$\theta_0 - \theta_\star = \frac{\delta}{\sqrt{N}}$ implies
$\zeta_N(\theta_0) \sim \mathcal{N}\left(-\mathbf{I}(\theta_\star)\,\delta\,,\, \mathbf{I}(\theta_\star)\right)$	$\zeta_N(\theta_0) \sim \mathcal{N}\left(-\mathbf{J}(\theta_\star)\,\delta\,,\, \mathbf{R}(\theta_\star)\right)$
$\mathbf{I}(\theta_\star) = -\frac{\partial}{\partial\theta} h(\theta, \theta_\star)\big\vert_{\theta=\theta_\star}$	$\mathbf{J}(\theta_\star) = -\frac{\partial}{\partial\theta} h(\theta, \theta_\star)\big\vert_{\theta=\theta_\star}$
$=$ Fisher information matrix	$\mathbf{R}(\theta_\star) = \lim_{N\to\infty} \mathrm{cov}\zeta_N(\theta_\star)$
$\chi(\theta_0) = \zeta_N^T(\theta_0)\mathbf{I}^{-1}(\theta_0)\zeta_N(\theta_0)$	$\chi(\theta_0) = \zeta_N^T(\theta_0)\mathbf{K}^{-1}\zeta_N(\theta_0)$
	$\mathbf{K}^{-1} =_{\rm def}\ \mathbf{R}^{-1}\mathbf{J}\left(\mathbf{J}^T\mathbf{R}^{-1}\mathbf{J}\right)^{-1}\mathbf{J}^T\mathbf{R}^{-1}$

Table 2: *The local approach, for the i.i.d. case with likelihood methods, and for the general case. In the last two lines,* $\mathbf{J}, \mathbf{R}, \mathbf{K}$*, should read* $\mathbf{J}(\theta_0)$*, etc.*

Eqn. (26) relates the identification error $\widehat{\theta}_N - \theta_\star$ to the efficient score $\zeta_N(\theta_\star)$ at the true model, under the local approach. Informally, Eqn. (26) relates identification to model validation. On the other hand, the asymptotic distribution of $\zeta_N(\theta_\star)$ is known to us: take $\delta = 0$ in (23). Combining this observation with (26), we re-derive the known result that the maximum likelihood estimator has asymptotic convergence rate

$$\boxed{\sqrt{N}\,\left(\widehat{\theta}_N - \theta_\star\right)\ \approx\ \mathcal{N}\left(0\,,\, \mathbf{J}^{-1}(\theta_\star)\mathbf{R}(\theta_\star)\mathbf{J}^{-T}(\theta_\star)\right)} \tag{27}$$

The comparison between i.i.d. case with likelihood methods and the general case of M-estimators is shown in Table 2.

5 Some use cases for the local approach

We discuss in this section some typical use cases of the local approach in the context of vibration mechanics.

5.1 Estimating confidence bounds in identification

As we have seen, M-estimators are appropriate for system identification in this context. Now, getting an estimate for the uncertainty of M-estimators can be difficult—for example, convergence rates for subspace methods are extremely complicated to get and impractical to compute Bauer et al. (1999). Of course, empirical estimates for such confidence bounds, *e.g.*, obtained from bootstrap or jacknife methods Davison and Hinkley (1997), can always be computed but at a prohibitive cost in computing time and amount of data.

By using the results of Section 4, we see that an alternative approach is possible. Assume we can compute (an estimate of) the Jacobian matrix $\mathbf{J}(\widehat{\theta}_N)$—typically, this is easier than computing the full convergence rate as given in (27). Then, by (23) applied with $\delta = 0$, we can estimate the missing matrix $\mathbf{R}(\widehat{\theta}_N)$ via an empirical covariance estimator of the "pseudo-score" $\zeta_N(\widehat{\theta}_N)$. And then we can conclude by (27).

This is illustrated on Fig. 4, for the case of a *mode shape*; we show, in red,

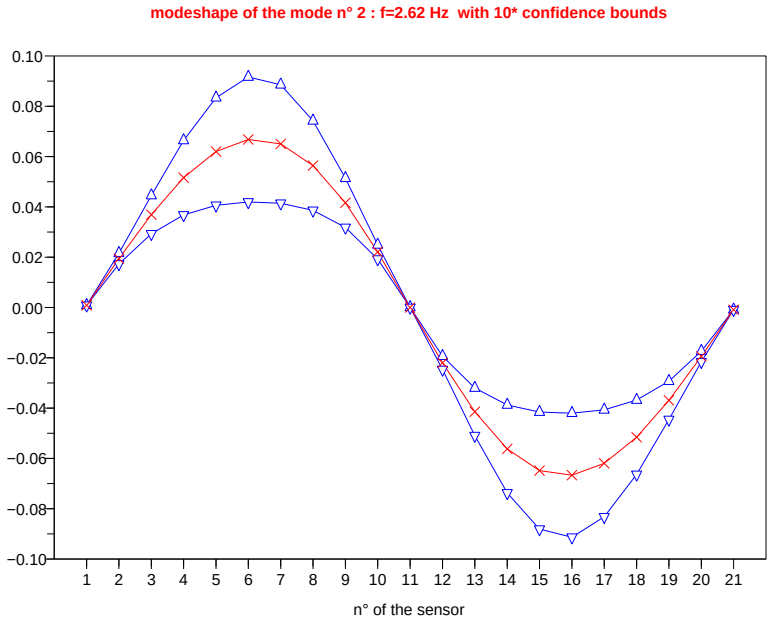

Figure 4: *Simulation example of Fig. 1. Confidence bounds on a mode shape (real part). To make it readable, the size of the bound is multiplied by a factor of* 10.

the considered mode shape (it is a subset of the coordinates of parameter θ), and its upper and lower bounds, in blue.

5.2 Model validation

Assume a nominal model θ_0 is at hand—it may have been identified on some data set. Then, assume that validation data are available. If the nominal model is close enough to the true one so that the local approach makes sense, then 1/ compute the pseudo-score $\zeta_N(\widehat{\theta}_N)$ and 2/ test whether the hypothesis that $\delta = 0$ in (23) can be accepted—this amounts to computing the χ^2-test of Table 2. Again, this requires computing estimates of the Jacobian matrix $\mathbf{J}(\widehat{\theta}_N)$ and the covariance matrix $\mathbf{R}(\widehat{\theta}_N)$, either based on identification data or on validation data. Fig. 2 illustrates the above method; see also Mevel and Goursat (2006) for further results, both on simulation and real data.

5.3 Change detection

In fact, the same approach allows us to perform damage detection: take for θ_0 a model of the safe structure and decide, based on current measurement data sets, whether the hypothesis that $\delta = 0$ in (23) can still be accepted. This amounts to comparing chi-square statistics of Table 2 to an appropriate threshold.

This is illustrated on Fig. 5 for the simulation example of Fig. 1. The two

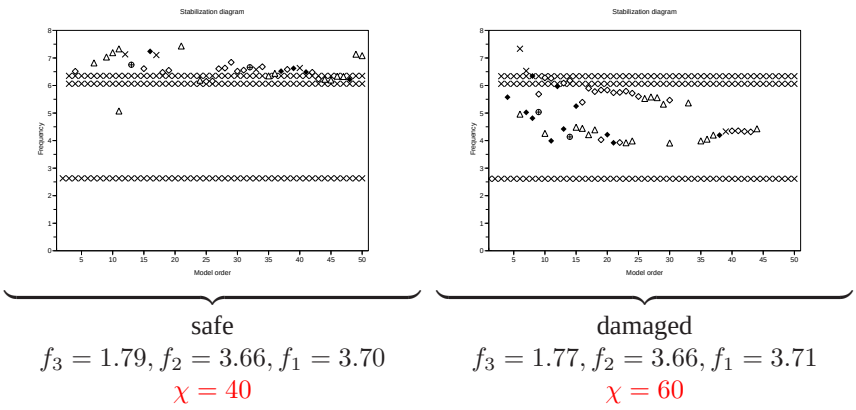

Figure 5: *Damage detection. Simulation example of Fig. 1. Stabilization diagram and corresponding chi-square tests, following Table 2.*

diagrams shown are (zooms of) *stabilization diagrams, i.e.*, plots of frequencies obtained from identifications performed with different model orders. In Fig. 5,

the estimates $f_i, i = 1, 2, 3$ for the frequencies were obtained via such a processing. Note that the change is adequately revealed by the chi-square statistics, although it causes a very small shift in the frequencies.

A real example. Fig. 6 shows a real example, using real data from the Z24

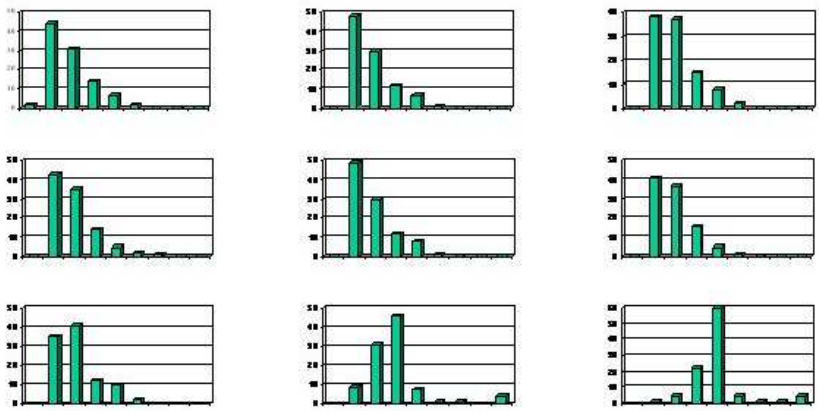

Figure 6: *Damage detection. The Z24 bridge: histograms of the damage detection test for the* 9 *month period. Note the shift toward larger values for the test, for the last two histograms, clearly reflecting the induced damages.*

bridge Mevel et al. (2003), tested in the framework of the BRITE-EURAM project SIMCES. The bridge used for validation is bridge Z24 in Canton Bern, Switzerland, connecting Koppigen and Utzenstorf. The bridge is a highway overpass of the A1, linking Bern and Zürich. Z24 is a prestressed bridge, with three spans, two lanes and 60 m overall length. The experiment was performed during year 1998. On 10.08.98 the settlement of pier was cut by 20 mm, and then 40mm on 12.08.98, 80mm on 17.08.98, 95mm on 18.08.98. Finally, a tilt was performed on the foundation on 19.08.98. Data were recorded, from January 1998 to September 1998. They were provided blindly to the competitors of the benchmark test. Results are shown on Fig. 6, in the form of 9 histograms, ordered as follows: first line: january, february, march; second line: april, etc. Each histogram depicts the distribution of the chi-square test computed for a number of records in the considered month. Histograms are close to identical for months 1–7. A change is visible on month 8 (the histogram is shifted toward larger values for the test), where damage was applied, and even more on month 9 since the structure was damaged for the entire month.

On-line monitoring. So far all algorithms we have discussed are off-line, since the (pseudo)score assumes a N-size data sample. A simple on-line form exists for the particular case in which the parameter θ for monitoring is a scalar—this does not require that the model itself shall be parameterized by a scalar, as the next sections show. For this case, the pseudo-score $\zeta_N(\theta_0) = \frac{1}{\sqrt{N}} \sum_1^N H(\theta_0, Y_k)$ is a scalar, and on-line monitoring with the local approach boils down to detecting a change in the mean of the sequence of random variables $H(\theta_0, Y_k)$. By the CLT, we can even do as if these variables were Gaussian and independent Benveniste et al. (1990). Very simple and effective CUSUM tests allow for detecting, *e.g.*, a change, from $\theta > 0$, to $\theta < 0$, see for example Mevel et al. (2005). For more general situations where the parameter for monitoring has to be a vector, no such simple solution exist, but of course the situation is not worse than for identification and on-line forms can be found on a case by case basis.

5.4 Focusing monitoring on some subspace

In some cases, one is not interested in monitoring θ for arbitrary changes. For example, the designer wishes to monitor one mode of her vibrating structure (by a mode, we mean the collection consisting of the frequency, damping, and mode shape or eigenvector). This can be captured in the following way. We assume that the monitoring is restricted to a subspace of the parameter space. This subspace is spaned by some parameter μ, of dimension smaller than that of θ. Thus, referring to eqn. (2), we have

$$\delta = L\mu, \ \dim(\mu) < \dim(\theta)$$

where we recall that δ is the normalized deviation between true and nominal model. The same substitution applies to (23), which gives raise to the appropriate modification of the chi-square statistics of Table 2, namely:

$$\chi_N^L(\theta_0) \quad =_{\text{def}} \quad \zeta_N^T(\theta_0)\, L \left(L^T \, \widehat{\mathbf{K}}_N(\theta_0)\, L \right)^{-1} L^T \, \zeta_N(\theta_0). \qquad (28)$$

This method is fine to decide between the cases $\mu = 0$ (no change) and $\mu \neq 0$ (some change occured for μ). But the analysis does not tell us how the above chi-square statistics reacts if indeed a change occurs that is not explicated by μ, *i.e.*, $\delta \notin \text{range}(L)$. For example, if μ collects a subset of the coordinates of δ, then it may be that a change affects other coordinates than those collected in μ. In general, unfortunately, the chi-square statistics will react to such a change. The solution to this problem consists in considering changes in other coordinates as a nuisance and rejecting them.

5.5 Sensitivity versus robustness: rejection of nuisances

Sensitivity of estimators or decision tests is of course an issue of primary interest. However, in many practical cases, the actual probability distribution of our observations may also depend on other parameters that are of no interest or are irrelevant for the problem under consideration. A typical example is the effect of temperature on a structure, which can possibly mask the onset of a damage, cf. Section 2.2.2. Such undesirable or irrelevant parameters are called *nuisances.*

In general, this is a difficult problem. However, if the effect of the nuisance on the actual distribution is known, then robustness of identification or tests can be predicted. More precisely, we now assume that our parameter θ decomposes as follows:

$$\theta \;=\; \left[\begin{array}{c} \theta^u \\ \theta^n \end{array}\right] \tag{29}$$

where superscripts ^{u} and ^{n} refer to the "useful" and "nuisance" parts of parameter θ, respectively. Following the decomposition (29) of θ, we decompose

$$\mathbf{J}^{-1}(\theta_\star)\mathbf{R}(\theta_\star)\mathbf{J}^{-T}(\theta_\star) \;=_{\rm def}\; \Delta(\theta_\star) \;=\; \left[\begin{array}{cc} \Delta^{uu}(\theta_\star) & \Delta^{un}(\theta_\star) \\ \Delta^{nu}(\theta_\star) & \Delta^{nn}(\theta_\star) \end{array}\right] \tag{30}$$

Assume that decomposition (30) yields a block-diagonal matrix. Then, by (27), estimates $\widehat{\theta}^u_N$ and $\widehat{\theta}^n_N$ are asymptotically independent, *i.e.,* they do not influence each other. Similarly, the pseudo-score (23) decomposes into two independent random vectors. This means that we can safely ignore the nuisance parameter and do as if we had only the reduced size parameter θ^u of interest in parameterizing the distribution as p_{θ^u}. Practically, this allows us to consider reduced size parameter identification by ignoring the nuisance in our M-estimator. This also allows us to perform testing for changes by using the same reduced size model: changes in the nuisance will not affect our test.

The closer to block-diagonal decomposition (30) is, the better is the situation. If this desirable situation fails to happen, then the residual

$$\zeta_N(\theta^u/\theta^n) \;=_{\rm def}\; \zeta^u_N(\theta_0) - \mathbf{E}_{\theta_\star}(\zeta^u_N(\theta_0) \mid \zeta^n_N(\theta_0)), \tag{31}$$

where $\mathbf{E}_{\theta_\star}(. \mid .)$ denotes conditional expectation under distribution $p_{\theta_\star}$, projects away the effect of the nuisance on the pseudo-score and provides a statistics that is robust against changes in the nuisance and still maximally sensitive to changes in the parameter of interest. Recall that the random vectors involved in (31) are asymptotically Gaussian for large data sets, which makes it easy to compute this conditional expectation.

We now indicate how to compute $\zeta_N(\theta^u/\theta^n)$ in practice. Let $H(\theta^u, \theta^n, Y_k)$ now depend both on the parameter θ^u of interest (e.g., referring to Section 2.2.2,

modes and mode shapes), and on the nuisance parameter θ^n (e.g., referring to Section 2.2.2, the temperature). Let

$$\zeta_N(\theta^u, \theta^n) \quad =_{\text{def}} \quad \frac{1}{\sqrt{N}} \sum_1^N H(\theta^u, \theta^n, Y_k)$$

be the pseudo-score as before. Then we can compute $\zeta_N(\theta^u/\theta^n)$ as follows:

$$\zeta_N(\theta^u/\theta^n) \quad =_{\text{def}} \quad U^T \zeta_N(\theta^u, \theta^n)$$

where matrix U is orthogonal, and of maximal rank such that

$$U^T \left(\frac{\partial}{\partial \theta^n} h(\theta^u, \theta^n; \theta^u_\star, \theta^n_\star) \Big|_{(\theta^u, \theta^n) = (\theta^u_\star, \theta^n_\star)} \right) \quad = \quad 0$$

The robust pseudo-score $\zeta_N(\theta^u/\theta^n)$ will not react to changes in θ^n and will do its best at reacting to changes in θ^u. Denote by

$$\overline{\chi}^u_N(\theta_0) \tag{32}$$

the resulting chi-square statistics.

Rejecting the temperature effect for the example of Section 2.2.2 is illustrated in Fig. 7, see Balmès, Basseville, Mevel and Nasser (2006). In this example, we show in blue (resp. in red) the statistics computed on the safe (resp. damaged) structure, for increasing temperature, with and without using the nuisance rejection mechanism. Note the local nature of the rejection mechanism: it is efficient for a small temperature change, but its effect decreases when the change gets larger, as revealed by the quadratic shape of the nuisance rejection test under no damage. In fact, a deeper use of the physical model involving temperature effects allows getting better rejection, see Fig. 8 and Basseville et al. (2006).

5.6 Diagnosis

When a change is detected by the chi-square statistics of Table 2, we would like to know which coordinates of parameter θ are most affected by this change. To this end, list the candidate subspaces for monitoring; for example, we could consider the case where all 1-dimensional subspaces $L_i, i = 1, \ldots, \dim(\theta)$ of the parameter space are for monitoring. Then, two different ways can be used:

- Compute the statistics $\chi^{L_i}_N(\theta_0), i = 1, \ldots, \dim(\theta)$ corresponding to (28), and compare them.
- Alternatively, compute the statistics $\chi^{u_i}_N(\theta_0), i = 1, \ldots, \dim(\theta)$ corresponding to (32), and compare them.

Many variants can be considered.

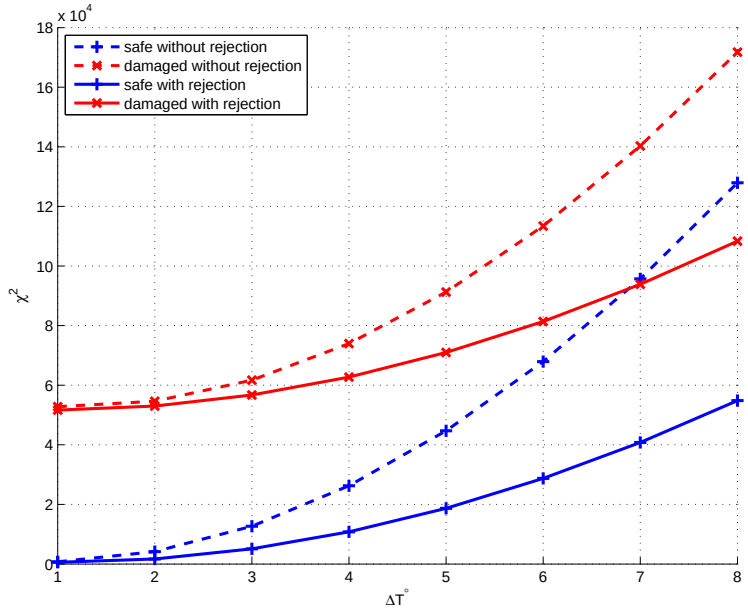

Figure 7: *Simulation example of Fig. 1, when temperature changes. Rejecting the temperature effect with the approach of Section 5.5. The figure shows that the robust chi-square test (in solid line) separates well safe from damaged states for the structure, for temperature ranging from 0 to 8 degrees Celsius. Note that the basic chi-square test (dashed lines) performs poorly, as expected from the frequency plots shown in Fig. 3.*

5.7 Diagnosis in terms of a design model

Now, the most interesting situation is when diagnosis is to be performed in terms of a *design model*, not in terms of the parameter space used in identification. Typically, design models are much richer than those used for identification. Therefore, they are of much larger dimension. Referring to the simulation example of Fig. 1, the dimension of the true system parameter space is nearly 10^5, whereas actual identification is performed with parameter spaces of dimension a few hundreds.

The approach we advocate is illustrated in Fig. 9. The key idea is to avoid mapping back changes, from the parameter space used for identification, to the space used for the design model. Reason for this is that such an inverse mapping is very ill conditioned—it corresponds to the difficult problem of *model updating*, *i.e.*, the problem of adjusting a high dimensional design model to best match a given identified one.

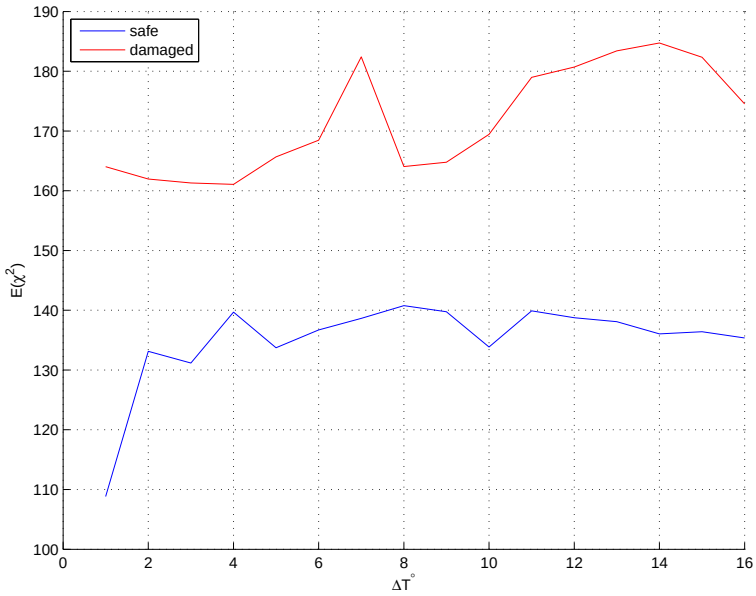

Figure 8: *Simulation example of Fig. 1, when temperature changes. Rejecting the temperature effect on a larger temperature range (*0–16 *degrees Celsius), by making a deeper (non local) use of the physical model.*

Instead, we use a direct mapping, from the physical space down to the modal space. This is well conditioned and requires computing Jacobians. We express the failure hypotheses of interest in the physical space. Their number would typically be very large (one failure for each parameter involved in each finite element). Mapping these failure hypotheses to the modal space yields a collection of points. Equip the modal space with the metrics associated to the local approach, namely the one defined by the matrix $\mathbf{K}^{-1}$ of Table 2. With this metrics, some failure hypotheses, when mapped to the modal space, become indistinguishable from the safe situation, they are not diagnosable. Some diagnosable failures cannot be isolated from each other. To take this into account, we perform clustering of the hypotheses in modal space, with the metrics associated to the local approach. This yields *macro-failure* hypotheses in the form of clustered classes $C_j, j \in J$. For each class C_j, we select a representative F_j. This representative spans a subspace in the modal space, we represent it by its column matrix L_j. Having done this, we are back to the diagnosis situation of Section 5.6.

This technique is illustrated on the bridge example of Fig. 1, see Balmès, Basseville, Mevel, Nasser and Zhou (2006). Fig. 10 shows the result of the clus-

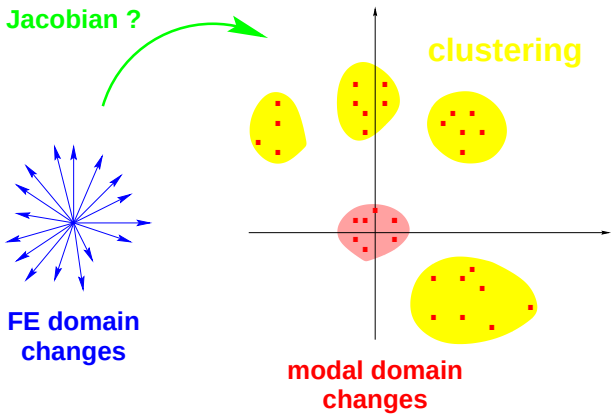

Figure 9: *Diagnosis in terms a design model: principle of the approach for the case of structural diagnosis with Finite Element (FE) parameters,* e.g., *Young modulus and mass of each element.*

tering. Fig. 11 shows the simulated damage along the x-axis, in red. The actual

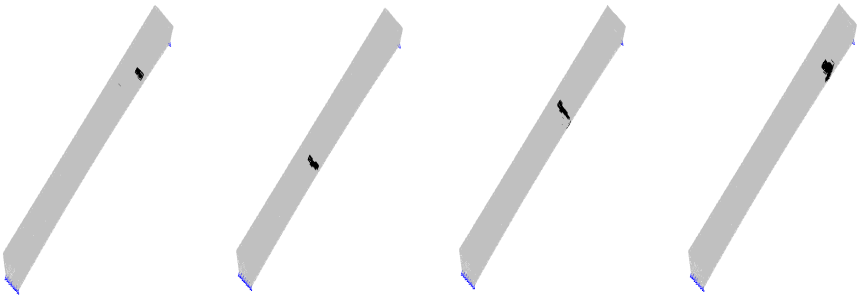

Figure 10: *FE Diagnosis. Showing the results of clustering. Four macro-failures are shown by marking in black the corresponding elements.*

localization of the damage, on the (y, z)-plane is shown on Fig. 12, top-right picture.

The results of diagnosis are shown on the other diagrams of Fig. 12. The very noisy plot on top-left displays the value of the chi-square test designed to monitor each hypothesis attached to each individual finite element, by using the focusing method of Section 5.4. On the second row, the plot on the left displays, for each finite element, the value of the chi-square test for the macro-failure containing the considered element (by using again the focusing method of Section 5.4). By

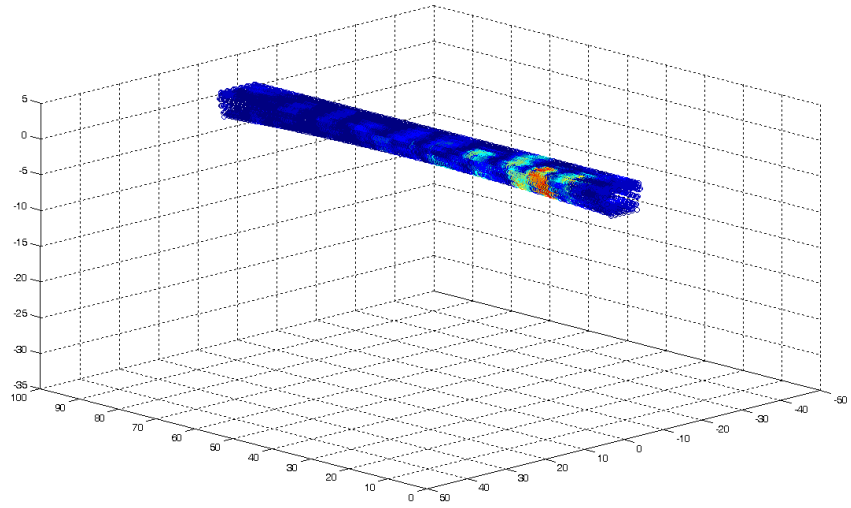

Figure 11: *FE Diagnosis. Showing the simulated damage, in red.*

construction, all elements belonging to the same macro-failure get the same value for the chi-square test.

The macro-failure that reacts most to this test is shown on the picture sitting on the right hand side—compare with the localization of the damage in Fig. 11 and Fig. 12, top right. On the third row we show the same, but for a different tuning of the clustering, where more macro-failures are found having fewer finite elements each. Note the similarity between the three plots displaying the values of the test for each case. Still, thanks to the clustering, the interpretation of the second (and third) plots is made much easier.

6 Concluding remarks: is the local approach a panacea?

Le Cam's local approach, combined with Ljung's clean separation between model structure and statistical considerations, allowed us to derive a general and powerful technique for confronting models to data. Having this, we were able to answer a large variety of questions involving the triple

{physical design model, identified model, data}

Well, all this sounds miraculous, and is indeed often so. Unfortunately, real life is not a dream. We list below a number of issues that, when not properly noticed,

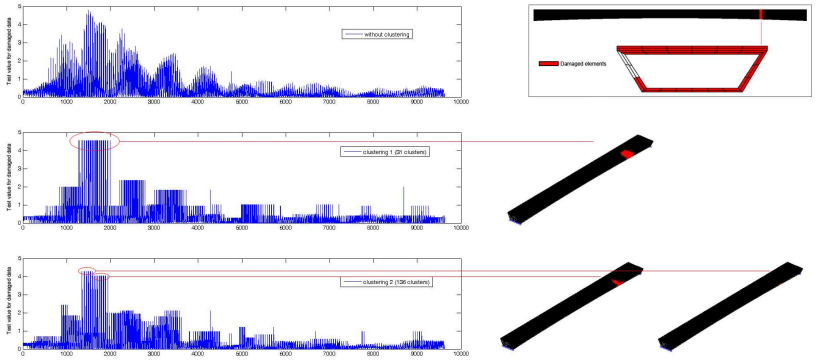

Figure 12: *FE Diagnosis. Showing the results.*

may prevent the local approach from delivering its promises:

What does the local hypothesis (2) mean in practice? What does it mean, for the nominal model, to sit "at an order of $1/\sqrt{N}$" away from the true parameter? Well, this is not really a concern when deducing convergence rates for identification algorithms, or estimating associated error covariance matrices: by definition, identification algorithms should provide "close-to-true" estimates.

In contrast, this becomes a concern when model validation or damage detection are considered: we actually do not know whether the deviation from true model, or actual damage, can be considered "small of order $1/\sqrt{N}$". The test statistics obtained from the local approach can always be considered and computed, but its efficiency may become questionable when changes become large!

The work Mevel and Goursat (2006) illustrates the above discussion. The experiments reported there show that the validation statistics behaves as a convex quadratic function (as expected) when the nominal and true values are close to each other, but this holds only locally.

All our analysis so far assumes that the true model belongs to the model set, *i.e.*, the actual distribution p exactly equals p_θ for some "true" value $\theta = \theta_\star$ for the parameter. However, in many practical situations, model reduction is enforced, which prevents this ideal situation from occurring.

This situation has been widely discussed by L. Ljung in his book Ljung (1999) on system idenfication. There is typically a "best model" in the model set, represented by parameter $\theta_\star$; in addition, a *true* parameter $\Theta_\star$ must be consid-

ered, typically in a space of larger dimension. The discrepancy $\delta\theta_\star = \theta_\star - \Theta_\star$ due to model reduction causes an additional bias to occur in the larger dimensional pseudo-score, *i.e.*, even if $\theta_0 = \theta_\star$, the pseudo-score is not zero-mean—compare with (20).

Worse, this bias may be conflicting with the local hypothesis (2): strictly speaking, we need that the bias on the pseudo-score caused by model reduction is of same magnitude order as the variance contribution due to noises and uncertainties.

This analysis provides theoretical guidelines for selecting model order: bias and variance should be of same magnitude order—recall the latter relates to sample size. Unfortunately, this guideline is not very easy to use in practice. See Zhang et al. (1994) for details on this subject.

The sensitivity and covariance matrices $\mathbf{J}(\theta_\star)$ and $\mathbf{R}(\theta_\star)$ of Section 4 may be difficult estimating, due to numerical instabilities. From our experience, most difficult is the estimation of $\mathbf{R}(\theta_\star)$. Of course, a poor estimate would impair the quality of confidence bounds in parameter identification, or the actual value of the test for model validation or damage detection.

A Appendix: selected modal analysis methods for use with the local approach

Here we cast some popular approaches to structural analysis in the framework of Section 4. We leave as an exercise for the reader to perform this task for her favorite algorithm.

Max Likelihood, prediction error, equation error, and related methods. Maximum likelihood approaches have been recently extensively considered for structural identification, both in time- and frequency-domain Verboven et al. (2004); Guillaume (2006).

Since the exact likelihood can be difficult expressing, several methods have been proposed that replace the exact likelihood for maximization by approximations of it. In the statistics litterature, these are known as *pseudo-likelihood* methods, whereas in the control litterature they are referred to as *prediction error* or *equation error* methods Ljung (1999). These methods generally consist in defining an error

$$e_k(\theta) = x_k - f(\theta, \mathcal{X}_{k-1}, \mathcal{U}_{k-1}),$$

where θ is the model parameter vector, u_k and x_k are the input and output at instant k, respectively, $\mathcal{U}_{k-1} = (u_{k-1}, \dots, u_{k-p})$ and similarly for $\mathcal{X}_{k-1}$, and $f(\theta, \mathcal{X}_{k-1}, \mathcal{U}_{k-1})$ is a "prediction" of x_k. Typical criteria for minimization are

then $\sum_{k=1}^{N} |e_k(\theta)|^2$, or its counterpart in the frequency domain. The function $H(\theta, y)$ is then taken equal to

$$\begin{array}{rcl} H(\theta, Y_k) & =_{\rm def} & \dfrac{\partial}{\partial\theta}|e_k(\theta)|^2, \text{ where} \\ e_k(\theta) & = & x_k - f(\theta, \mathcal{X}_{k-1}, \mathcal{U}_{k-1}) \text{ and} \\ Y_k & =_{\rm def} & (x_k, \mathcal{X}_{k-1}, \mathcal{U}_{k-1}). \end{array}$$

Of course, it is important that $\frac{\partial}{\partial\theta} e_k(\theta)$ can be effectively computed, at a reasonable cost. From our previous analysis, we see that dealing with uncertainties require estimating $\mathbf{R}(\theta)$ and $\mathbf{J}(\theta)$; we discussed this in Section 6.

Subspace methods for output-only structural identification. Consider the following linear system

$$\left\{ \begin{array}{rcccl} x_k & = & Ax_{k-1} & + & v_k \\ y_k & = & Cx_{k-1} & + & w_k \end{array} \right. \tag{33}$$

where $k \in \mathbb{Z}$, x is the $\mathbb{R}^p$-valued state, v and w are unobserved input disturbances, and y is the $\mathbb{R}^q$-valued observed output. The problem we consider is the *identification of the pair* (C, A) *up to a change of basis in the state space of system* (33). Equivalently, we identify the pairs $(\lambda, C\varphi_\lambda)$, where λ ranges over the set of eigenvalues of A (the poles of system (33)) and φ_λ are the corresponding eigenvectors. Said in words, we consider the problem of *eigenstructure identification.*

The *subspace method* for eigenstructure identification can be seen as follows. Consider the following parameter vector θ:

$$\theta =_{\rm def} \begin{pmatrix} \Lambda \\ \text{vec}\,\Phi \end{pmatrix}$$

where Λ is the diagonal matrix whose elements are the eigenvalues λ and Φ is the matrix whose columns are the φ_λ's defined above. Then, consider the following observability matrix in modal form:

$$\mathcal{O}(\theta) =_{\rm def} \begin{bmatrix} \Phi \\ \Phi\Lambda \\ \Phi\Lambda^2 \\ \vdots \\ \Phi\Lambda^m \end{bmatrix},$$

and let $\mathcal{S}(\theta)$ be a matrix of maximal rank such that

$$\mathcal{S}^T(\theta)\mathcal{O}(\theta) = 0.$$

Assume a N-size sample $y_{-m}, \ldots, y_{N+m}$. Then, define Y_k and $\check{Y}_k$ by $Y_k^T =_{\rm def} \left[y_k^T \; y_{k-1}^T \; \ldots \; y_{k-m}^T\right]$ and $\check{Y}_k^T =_{\rm def} \left[y_k^T \; y_{k+1}^T \; \ldots \; y_{k+m}^T\right]$, respectively, and consider the empirical *Hankel matrix*:

$$\mathcal{H} =_{\rm def} \sum_{k=0}^{N} \check{Y}_k Y_k^T = \begin{bmatrix} \hat{R}_0 & \hat{R}_1 & \hat{R}_2 & \ldots & \hat{R}_m \\ \hat{R}_1 & \hat{R}_2 & \ldots & \ldots & \hat{R}_{m+1} \\ \hat{R}_2 & \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ \hat{R}_m & \hat{R}_{m+1} & \ldots & \ldots & \hat{R}_{2m} \end{bmatrix}$$

where $\hat{R}_i$ is the ith empirical covariance matrix of observation y_k. Then, the subspace algorithm for identifying the eigenstructure θ is given by:[3]

$$\widehat{\theta}_N =_{\rm def} \arg_\theta \left[\, \zeta_N(\theta) \; = \; 0 \right]$$

where

$$\begin{aligned} \zeta_N(\theta) \quad &=_{\rm def} \quad \frac{1}{\sqrt{N}} \mathcal{S}^T(\theta)\mathcal{H} \; = \; \frac{1}{\sqrt{N}} \sum_{k=0}^{N} \mathcal{S}^T(\theta) \check{Y}_k Y_k^T \\ &=_{\rm def} \quad \frac{1}{\sqrt{N}} \sum_{k=1}^{N} H(\theta, Y_k, \check{Y}_k)\,, \end{aligned}$$

which shows that subspace methods are just an instance of M-estimators. Hence Section 4 applies. This is the method we have used in the reported examples.

References

Balmès, E., Basseville, M., Mevel, L. and Nasser, H. (2006). Handling the temperature effect in vibration-based monitoring of civil structures: a combined subspace-based and nuisance rejection approach, *Proc. of the 6th Symposium on Fault Detection, Supervision and Safety of Technical Processes (*SAFEPROCESS*)*, Beijing, China.

Balmès, E., Basseville, M., Mevel, L., Nasser, H. and Zhou, W. (2006). Statistical model-based damage localization: a combined subspace-based and substructuring approach, *Proc. of the 4th IASC World Conference on Structural Control and Monitoring*, San Diego, CA.

Basseville, M., Bourquin, F., Mevel, L., Nasser, H. and Treyssède, F. (2006). Handling the temperature effect in SHM: combining a subspace-based statistical test and a temperature-adjusted null space, *Proceedings of the 3rd European Workshop on Structural Health Monitoring*, Granada, Spain.

[3]The reader may be surprised by this claim. She certainly implements the subspace method by SVD factorizing the Hankel matrix, and then computing the pair (C, A) from the left factor. These two formulations are indeed equivalent, up to asymptotically neglectible terms. The interest of the formulation used here is that it casts subspace methods in the general framework of M-estimators.

Basseville, M. and Nikiforov, I. V. (1993). *Detection of Abrupt Changes – Theory and Application*, Prentice-Hall. Book downloadable from the page `http://www.irisa.fr/sisthem/kniga/`.

Bauer, D., Deistler, M. and Scherrer, W. (1999). Consistency and asymptotic normality of some subspace algorithms for systems without observed inputs, *Automatica* **35**(7): 1243–1254.

Benveniste, A., Basseville, M. and Mousatakides, G. (1987). The asymptotic local approach to change detection and model validation, *IEEE Trans. Automat. Contr.* **AC-32**(7): 583–592.

Benveniste, A. and Fuchs, J.-J. (1985). Single sample modal identification of a nonstationay stochastic process, *IEEE Trans. Automat. Contr.* **AC-30**: 66–74.

Benveniste, A., Métivier, M. and Priouret, P. (1990). *Adaptive Algorithms and Stochastic Approximations*, Springer-Verlag, New York.

Brown, R., Durbin, J. and Evans, J. (1975). Techniques for testing the constancy of regression relationships over time., *Jal. of the Royal statistical Society (B)* **2**: 149–192.

Davies, R. B. (1973). Asymptotic inference in stationary Gaussian time-series, *Adv. Appl. Prob.* **5**: 469–497.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap methods and their application*, Cambridge series in statistical and probabilistic mathematics, Cambridge University Press.

Delyon, B., Juditsky, A. and Benveniste, A. (1997). On the relationship between identification and local tests, *Technical Report No 1104*, Research Report IRISA.
URL: *http://www.irisa.fr/bibli/publi/pi/1997/1104/1104.html*

Gevers, M. (2005). Address in the honour of Lennart Ljung when Lennart received Honoris Causa degree in Katholieke University of Leuven, Belgium.

Guillaume, P. (2006). Multivariable frequency domain system identification algorithms for modal analysis, *Proc. of 14th IFAC Symposium on System Identification, SYSID'06*, Newcastle, Australia.

Hinkley, D. V. (1970). Inference about the change point in a sequence of random variables, *Biometrika* **57**(1): 1–17.

Huber, P. J. (1981). *Robust Statistics*, Wiley Series in Probability and Statistics, Wiley.

Le Cam, L. (1960). Locally asymptotic normal families of distributions, *University of California Publications in Statistics* **3**(2): 37–98.

Le Cam, L. (1986). *Asymptotic methods in statistical decisions theory*, Series in statistics, Springer, New York.

Lehmann, E. L. (1968). *Testing statistical Hypotheses*, 2nd edn, Wiley.

Ljung, L. (1977a). Analysis of recursive stochastic algorithms, *IEEE Trans. Automat. Contr.* **22**(8): 551–575.

Ljung, L. (1977b). On positive real transfer functions and the convergence of some recursive schemes., *IEEE Trans. Automat. Contr.* **22**(8): 539–551.

Ljung, L. (1978). Convergence analysis of parametric identification methods, *IEEE Trans. Automat. Contr.* **23**(5): 770–783.

Ljung, L. (1999). *System Identification – Theory For the User*, PTR Prentice Hall, Upper Saddle River, N.J.

Ljung, L. and Soderström, T. (1983). *Theory and practice of recursive identification*, MIT Press, Cambridge, Mass.

Mevel, L., Basseville, M. and Benveniste, A. (2005). Fast in-flight detection of flutter onset - A statistical approach, *AIAA Jal Guidance, Control, and Dynamics* **28**(3): 431–438.

Mevel, L. and Goursat, M. (2006). Model validation by using a damage detection test, *Proc. of the 24th International Modal Analysis Conference (IMAC-XXIV)*, Saint Louis, MI.

Mevel, L., Goursat, M. and Basseville, M. (2003). Stochastic subspace-based structural identification and damage detection and localization – Application to the Z24 bridge benchmark, *Mechanical Systems and Signal Processing, Special Issue on COST F3 benchmarks* **17**(1): 143–151.

Nikiforov, I. V. (1983). *Sequential detection of abrupt changes in time series properties*, Nauka, Moscow.

Roussas, G. G. (1972). *Contiguity of probability measures, some applications in statistics*, Cambridge University Press, New York.

Shiryaev, A. N. (1961). The problem of quickest detection of the destruction of a stationary regime, *Dokl. Acad. Nauk. SSSR* **138**(5): 1039–1042.

Verboven, P., Guillaume, P., Cauberghe, B., Vanlanduit, S. and Parloo, E. (2004). Modal parameter estimation from input/output fourier data using frequency-domain maximum likelihood identification, *Jal Sound and Vibration* **276**(3-5): 957–979.

Wald, A. (1947). *Sequential analysis*, Wiley, New York.

Willsky, A. S. and Jones, H. L. (1967). A generalized likelihood ratio approach to detection and estimation of jumps in linear systems, *IEEE Trans. Automat. Contr.* **21**(1): 108–112.

Zhang, Q., Basseville, M. and Benveniste, A. (1994). Early warning of slight changes in systems and plants with application to condition based maintenance, *Automatica* **30**(1): 95–114.