Fisher matrix methods for transfer function measurement

E. D. Hall

2015–10–17

Abstract: Many optical, mechanical, and optomechanical systems are assumed to be linear and time-invariant. Under this assumption, the parameters of such a system can be estimated by measuring the system's transfer function. In the case of systems such as alignment control systems, seismic isolation systems, and mechanical suspensions, transfer function measurement can be quite expensive, as it can require exciting the system down to millihertz frequencies.

With a little forethought, one can choose the measurement's excitation frequencies (and their amplitudes) so as to maximize the amount of information learned about the system's parameters. One method for quantifying the amount of information gained is to compute the Fisher matrix of the measurement. [1] The inverse of the Fisher matrix provides a lower bound on the covariance matrix of the estimated parameters.

1 General discussion

We define the following notation: we have an LTI system whose transfer function $H(f)$ depends on a certain set of parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_M)$. We can probe this system by exciting its input with a signal $x(f)$ and reading back the response $y(f) = H(f)x(f) + n(f)$, where $n(f)$ is some readout noise. Our goal is to produce an estimate $\hat{H}(\boldsymbol{\theta}; f)$ given our observed response $y(f)$, our excitation $x(f)$, and our estimate of the noise $n(f)$.

In practice, we excite the system with sinusoids at frequencies $f_1, f_2, \ldots, f_N$. We record the (in principle complex) excitation amplitudes $x_1, x_2, \ldots, x_N$ and the response amplitudes $y_1, y_2, \ldots, y_N$. These amplitudes have been corrupted by noise whose amplitudes are $n_1, n_2, \ldots, n_N$; in general, we have

$$y_\alpha = H_\alpha x_\alpha + n_\alpha; \qquad \alpha = 1, 2, \ldots, N, \tag{1}$$

where $H_\alpha = H(f_\alpha)$. On the other hand, given an estimate $\hat{H}(\boldsymbol{\theta}; f)$ of the system, we can write down a set of estimated amplitudes $\hat{y}_1, \hat{y}_2, \ldots, \hat{y}_N$, with

$$\hat{y}_\alpha = \hat{H}_\alpha(\boldsymbol{\theta}) x_\alpha; \qquad \alpha = 1, 2, \ldots, N, \tag{2}$$

Our goal is to find the value of $\boldsymbol{\theta}$ which makes the estimated responses $\{\hat{y}_\alpha\}$ approach the observed responses y_α. To this end, we can write down a likelihood function $\mathcal{L}(\boldsymbol{\theta}) \propto p(\{y_\alpha\}|\boldsymbol{\theta})$, where $p(\{y_\alpha\}|\boldsymbol{\theta})$ is the probability of having observed the amplitudes $\{y_\alpha\}$ given a certain value of $\boldsymbol{\theta}$. From here on we will assume that the noise is Gaussian, which results in a likelihood function[a]

$$\mathcal{L}(\boldsymbol{\theta}) \propto \exp\left[-\sum_{\alpha=0}^{N-1} \frac{|y_\alpha - \hat{y}_\alpha(\boldsymbol{\theta})|^2}{2|n_\alpha|^2}\right]. \tag{3}$$

How should we place our N frequencies so as to maximize the amount of information we can learn about H? Intuitively, we know we should choose our frequencies so as to maximize the curvature of $\mathcal{L}$ (or, equivalently, the curvature of $\ln\mathcal{L}$) with respect to $\boldsymbol{\theta}$. To find an expression for the curvature, we vary $\boldsymbol{\theta}$ and keep track of terms up to second order:

$$\ln\mathcal{L}(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \simeq \ln\mathcal{L}\big|_{\boldsymbol{\theta}} + \left.\frac{\partial[\ln\mathcal{L}]}{\partial\theta_i}\right|_{\boldsymbol{\theta}} \delta\theta_i + \frac{1}{2}\left.\frac{\partial^2[\ln\mathcal{L}]}{\partial\theta_i \partial\theta_j}\right|_{\boldsymbol{\theta}} \delta\theta_i\,\delta\theta_j, \tag{4}$$

where summation over i and j is understood. Once we have found parameters $\boldsymbol{\theta}_0$ which maximize $\ln\mathcal{L}$, the first-derivative terms will vanish, leaving only the second-derivative (i.e., curvature) terms. These curvature terms are elements of the Fisher matrix $\mathcal{F}$:[b]

$$\mathcal{F}_{ij} = -\left.\frac{\partial^2[\ln\mathcal{L}]}{\partial\theta_i \partial\theta_j}\right|_{\boldsymbol{\theta}_0} \tag{5a}$$

$$= \frac{\partial^2}{\partial\theta_i \partial\theta_j}\left[\sum_\alpha \frac{|y_\alpha - \hat{y}_\alpha(\boldsymbol{\theta})|^2}{2|n_\alpha|^2}\right]\Bigg|_{\boldsymbol{\theta}_0} \tag{5b}$$

$$= \sum_\alpha \frac{1}{2|n_\alpha|^2}\left[\frac{\partial\hat{y}_\alpha^*}{\partial\theta_i}\frac{\partial\hat{y}_\alpha}{\partial\theta_j} - \left(y_\alpha^* - \hat{y}_\alpha^*\right)\frac{\partial^2\hat{y}_\alpha}{\partial\theta_i\partial\theta_j} + \mathrm{cc}\right]\Bigg|_{\boldsymbol{\theta}_0}. \tag{5c}$$

If our estimate $\hat{y}$ is unbiased, we expect $y_\alpha - \hat{y}_\alpha \to 0$, and thus

$$\mathcal{F}_{ij} = \sum_\alpha \frac{1}{|n_\alpha|^2} \mathrm{Re}\left[\frac{\partial\hat{y}_\alpha^*}{\partial\theta_i}\frac{\partial\hat{y}_\alpha}{\partial\theta_j}\right]\Bigg|_{\boldsymbol{\theta}_0}. \tag{6}$$

Since $\hat{y}_\alpha = \hat{H}_\alpha x_\alpha$, we have

$$\mathcal{F}_{ij} = \sum_\alpha \frac{|x_\alpha|^2}{|n_\alpha|^2} \mathrm{Re}\left[\frac{\partial\hat{H}_\alpha^*}{\partial\theta_i}\frac{\partial\hat{H}_\alpha}{\partial\theta_j}\right]\Bigg|_{\boldsymbol{\theta}_0}. \tag{7}$$

which is the discrete-frequency equivalent of the expression found by L. Price. [1] Following the usual convention, we'll write $\sigma(H_\alpha) = |n_\alpha/x_\alpha|$.[c]

[a] If one is not comfortable with likelihood functions, note that $-\ln\mathcal{L}$ is equivalent to the usual χ^2 statistic used for parameter estimation.

[b] Note that $\mathcal{F} = -\mathcal{H}$, where $\mathcal{H}$ is the Hessian of $\ln\mathcal{L}$ evaluated at $\boldsymbol{\theta}_0$. Strictly speaking $\mathcal{F}$ is the *observed Fisher information*, which is to be contrasted with the *expected Fisher information*. See Efron and Hinkley [2] for more information.

[c] Alternatively, we can decompose each observation into real/imaginary parts: $y_\alpha = y_\alpha^{(\mathrm{r})} + \mathrm{i}y_\alpha^{(\mathrm{i})}$, and likewise for H_α and n_α. Then we can write down an alternative Fisher matrix

$$\mathcal{F}_{ij} = \sum_{\alpha=1}^{N}\sum_{\beta\in\{\mathrm{r},\mathrm{i}\}} \frac{|x_\alpha|^2}{\left|n_\alpha^{(\beta)}\right|^2}\frac{\partial\hat{H}_\alpha^{(\beta)}}{\partial\theta_i}\frac{\partial\hat{H}_\alpha^{(\beta)}}{\partial\theta_j}. \tag{8}$$

In the case $n_\alpha^{(\mathrm{r})} = n_\alpha^{(\mathrm{i})}$, this reduces to the definition given above.

The inverse of the Fisher matrix provides a lower bound for the covariance matrix:

$$\Sigma \geq \mathcal{F}^{-1}, \tag{9}$$

where the inequality is understood to be elementwise. This is the so-called Cramér–Rao bound.

2 Single-frequency estimation

As an example, we apply the above concepts to the problem of estimating the parameters of a single-pole system $H(f) = k/(1 + \mathrm{i}f/p)$ using an excitation at only one frequency. This system describes, for example, the optical response of a resonant Fabry–Pérot cavity to length or frequency perturbations. In this context, k is called the optical gain and p is called the cavity pole. For concreteness we can consider the aLIGO DARM cavity, for which (during O1) we have $k \simeq 3.2$ mA/pm and $p \simeq 350$ Hz.

For this system, our parameter vector is $\boldsymbol{\theta} = (k, p)$, so the Fisher matrix will be 2×2. As a start, we'll consider the case where we have only a single excitation at a frequency f_1. This results in only a single observation y_1.[d]

In this instance, the Fisher matrix is

$$\mathcal{F} = \frac{1}{\sigma_1^2} \begin{bmatrix} \dfrac{p^2}{p^2 + f_1^2} & \dfrac{kpf_1^2}{(p^2+f_1^2)^2} \\ \dfrac{kpf_1^2}{(p^2+f_1^2)^2} & \dfrac{k^2 f_1^2}{(p^2+f_1^2)^2} \end{bmatrix}, \tag{10}$$

and as a result, the covariance matrix Σ is bounded elementwise from below by $\mathcal{F}^{-1}$:

$$\Sigma \geq \sigma_1^2 \begin{bmatrix} \dfrac{1}{p^4}(p^2+f_1^2)^2 & -\dfrac{(p^2+f_1^2)^2}{kp^3} \\ -\dfrac{(p^2+f_1^2)^2}{kp^3} & \dfrac{(p^2+f_1^2)^3}{k^2p^2f_1^2} \end{bmatrix}. \tag{11}$$

From here, the goal is to choose f_1 so as to provide the "optimal" Σ. To make progress, we need to make assumptions about σ_1, which means making assumptions about the excitation amplitude x_1 and the readout noise amplitudes n_1. The next sections explore two simple cases for x_1 and n_1.

2 Flat excitation and white readout noise

To start with, we'll assume that the readout noise n_α is Gaussian and white as a function of frequency. (In the case of aLIGO DARM, the assumption of whiteness is true only above 100 Hz or so.) We'll also assume that the excitation amplitude is flat: $x_\alpha = x_0$ for all α. Then $\sigma_1 = \sigma(f_1) \equiv \sigma$ (i.e., it is independent of frequency). In this case, the elements of Σ are minimized as follows:

[d] However, since y_1 is complex, it contains two pieces of information: $y_1 = y_1^{(\mathrm{r})} + \mathrm{i}y_1^{(\mathrm{i})}$.

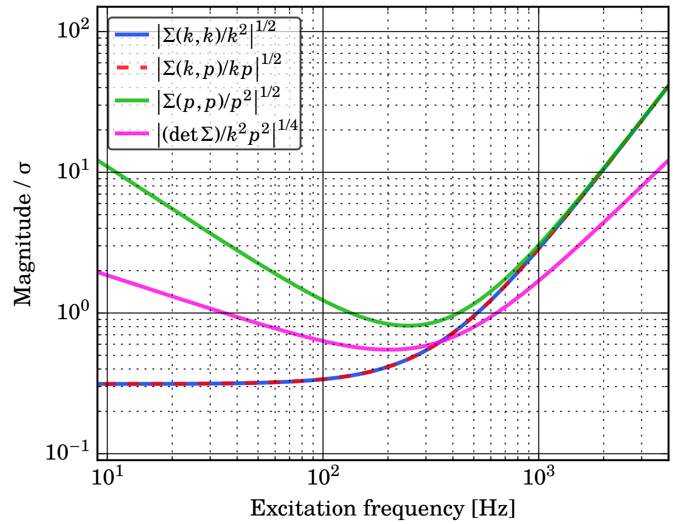

Figure 1: Normalized Cramér–Rao bounds on the covariance matrix Σ for a single-frequency TF estimate of the system $H(f) = k/(1 + \mathrm{i}f/p)$, with $k = 3.2$ mA/pm and $p = 350$ Hz. Here we have assumed a white readout noise and a flat excitation amplitude.

1. To minimize $\Sigma(k,k)$, one should choose $f_1 = 0$ Hz. This results in
$$\Sigma(0) \geq \begin{bmatrix} \sigma^2 & -\sigma^2 p/k \\ -\sigma^2 p/k & \infty \end{bmatrix}, \tag{12}$$
which is evidently unacceptable for simultaneous estimation of k and p.
2. To minimize $\Sigma(k,p)$, one should again choose $f_1 = 0$ Hz.
3. To minimize $\Sigma(p,p)$, one should choose $f_1 = p/\sqrt{2}$. This results in
$$\Sigma(p/\sqrt{2}) \geq \frac{9}{4} \begin{bmatrix} \sigma^2 & -\sigma^2 p/k \\ -\sigma^2 p/k & 3\sigma^2 p^2/k^2 \end{bmatrix}. \tag{13}$$
4. To minimize $\det\Sigma$, one should choose $f_1 = p/\sqrt{3}$. This results in
$$\Sigma(p/\sqrt{3}) \geq \frac{16}{9} \begin{bmatrix} \sigma^2 & -\sigma^2 p/k \\ -\sigma^2 p/k & 4\sigma^2 p^2/k^2 \end{bmatrix}. \tag{14}$$

In figure 1 we plot the elements of Σ, along with $\det\Sigma$, as a function of f_1, assuming that the Cramér–Rao bound is saturated.

2 Constant-SNR excitation

Often we can do better than a flat excitation amplitude. If we already have reasonable knowledge of k, p, and n, it is desirable to aim for the measurement to have a constant SNR. For our purposes, we define this as $\rho_\alpha = |y_\alpha/n_\alpha| = |x_\alpha H_\alpha/n_\alpha|$. Therefore, we choose $|x_\alpha| = \rho|n_\alpha|/|H_\alpha|$, where ρ is our target SNR, and hence $\sigma_\alpha^{(\mathrm{r,i})} = |H_\alpha|/\rho$.

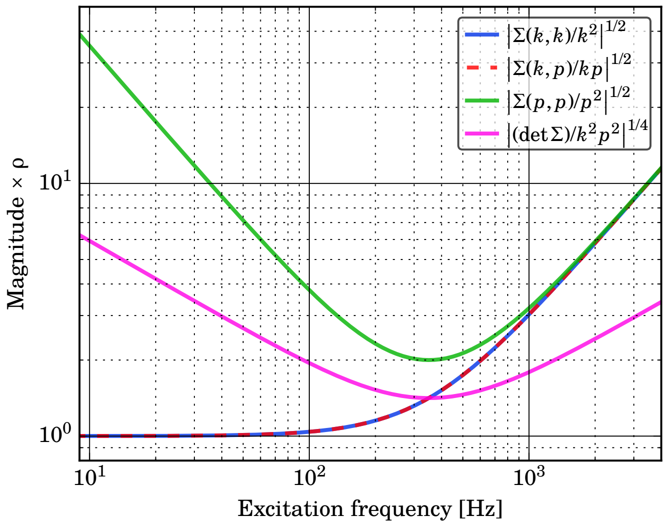

Figure 2: Normalized Cramér–Rao bounds on the covariance matrix Σ for a single-frequency TF estimate of the system $H(f) = k/(1+\mathrm{i}f/p)$, with $k = 3.2\,\mathrm{mA/pm}$ and $p = 350\,\mathrm{Hz}$. Here we have assumed that the response y has a constant SNR ρ above the readout noise n.

With this choice of excitation amplitude, the Fisher matrix is

$$\mathcal{F} = \rho^2 \begin{bmatrix} \frac{1}{k^2} & \frac{f_1^2}{kp\left(p^2+f_1^2\right)} \\ \frac{f_1^2}{kp\left(p^2+f_1^2\right)} & \frac{f_1^2}{p^2\left(p^2+f_1^2\right)} \end{bmatrix} \tag{15}$$

and the covariance matrix satisfies

$$\Sigma \geq \frac{1}{\rho^2} \begin{bmatrix} \frac{k^2}{p^2}\left(p^2+f_1^2\right) & -\frac{k}{p}\left(p^2+f_1^2\right) \\ -\frac{k}{p}\left(p^2+f_1^2\right) & \frac{1}{f_1^2}\left(p^2+f_1^2\right)^2 \end{bmatrix}. \tag{16}$$

As expected, the bound on the covariance goes down like ρ^2. To minimize the bounds on both $\Sigma(k,k)$ and $\Sigma(k,p)$, one should again choose $f_1 = 0\,\mathrm{Hz}$. However, this again results in $\Sigma(p,p) = \infty$:

$$\Sigma(0) \geq \begin{bmatrix} k^2/\rho^2 & -kp/\rho^2 \\ -kp/\rho^2 & \infty \end{bmatrix}. \tag{17}$$

This time, the bounds on both $\Sigma(p,p)$ and $\det\Sigma$ are minimized by choosing $f_1 = p$:

$$\Sigma(p) \geq 2 \begin{bmatrix} k^2/\rho^2 & -kp/\rho^2 \\ -kp/\rho^2 & 2p^2/\rho^2 \end{bmatrix}. \tag{18}$$

References

[1] Larry Price. Optimal experimental design: Introduction and application to system identification. Technical Report LIGO-G1400084-v1, LIGO Caltech, 2014. URL https://dcc.ligo.org/LIGO-G1400084.

[2] Bradley Efron and David V Hinkley. Assessing the accuracy of the maximum likelihood estimator: Observed versus expected fisher information. *Biometrika*, 65(3): 457–483, 1978.