Optimal Experimental Design: Utility Function for a Linear Model

May 29, 2017

1 Information Theory Measures

For a random variable X with distribution $p(x)$ the *information entropy* is defined as,

$$H(X) = -\int p(x)\,\log p(x)\,\mathrm{d}x\,. \tag{1}$$

The *mutual information* between the random variable X and a random variable Y with distribution $p(y)$ and joint distribution $p(x,y)$ is defined as,

$$I(X;Y) = \int\int p(x,y)\,\log\frac{p(x,y)}{p(x)p(y)}\,\mathrm{d}x\,\mathrm{d}y\,. \tag{2}$$

The mutual information can be identified as the relative entropy (Kullback-Leibler divergence) from the product distribution to the joint distribution,

$$I(X;Y) = D_{KL}\big(p(x,y)\,\|\,p(x)p(y)\big)\,. \tag{3}$$

Finally, if $X \sim \mathcal{N}(\mu,\Sigma)$ in $\mathbb{R}^M$, then

$$H(X) = \frac{1}{2}\log\Big((2\pi\mathrm{e})^M|\Sigma|\Big)\,. \tag{4}$$

2 Utility Function

Let ϑ be the parameters of a model and y the measurements obtained from an experiment with probability distributions $p(\vartheta)$ and $p(y)$ respectively. Moreover, let d be the design variable which will be treated as a parameter and not a random variable in this example. The *utility function* is defined as the mutual information of ϑ and y, with parameter d, is given by,

$$U(d) := I(\vartheta;y|d) = \int\int p(\vartheta,y|d)\,\log\frac{p(\vartheta,y|d)}{p(\vartheta|d)p(y|d)}\,\mathrm{d}\vartheta\,\mathrm{d}y\,. \tag{5}$$

Notice that the above quantity is not the conditional mutual information, d is treated as a parameter here. Under the assumption that the prior information of the parameters does not depend on the design, i.e.,

$$p(\vartheta|d) = p(\vartheta)\,, \tag{6}$$

we can write (5) as,

$$\begin{aligned}
U(d) &= \int\int p(\vartheta,y|d)\,\log\frac{p(\vartheta|y,d)p(y|d)}{p(\vartheta)p(y|d)}\,\mathrm{d}\vartheta\,\mathrm{d}y \\
&= \int\int p(\vartheta,y|d)\,\log\frac{p(\vartheta|y,d)}{p(\vartheta)}\,\mathrm{d}\vartheta\,\mathrm{d}y \\
&= \int\int p(\vartheta,y|d)\,\log p(\vartheta|y,d)\,\mathrm{d}\vartheta\,\mathrm{d}y - \int\int p(\vartheta,y|d)\,\log p(\vartheta)\,\mathrm{d}\vartheta\,\mathrm{d}y \\
&= \int\int p(\vartheta|y,d)\,p(y|d)\,\log p(\vartheta|y,d)\,\mathrm{d}\vartheta\,\mathrm{d}y - \int\left(\int p(\vartheta,y|d)\,\mathrm{d}y\right)\log p(\vartheta)\,\mathrm{d}\vartheta \\
&= \int\left(\int p(\vartheta|y,d)\,\log p(\vartheta|y,d)\,\mathrm{d}\vartheta\right)p(y|d)\,\mathrm{d}y - \int p(\vartheta|d)\,\log p(\vartheta)\,\mathrm{d}\vartheta \\
&= -\int H(\vartheta|y,d)\,p(y|d)\,\mathrm{d}y - \int p(\vartheta)\,\log p(\vartheta)\,\mathrm{d}\vartheta\,,
\end{aligned} \tag{7}$$

where we used the fact that

$$p(\vartheta, y|d) = p(\vartheta|y, d)p(y|d) \ , \tag{8}$$

and

$$\int p(\vartheta, y|d)\, \mathrm{d}y = p(\vartheta|d) \ . \tag{9}$$

Finally,

$$U(d) = H(\vartheta) - \mathbb{E}_{y|d}\big[H(\vartheta|y, d)\big] \ , \tag{10}$$

which is the information entropy of the prior minus the expectation under all possible measurements of the information entropy of the posterior.

3 Linear Model

Consider the linear model,

$$y = A\vartheta + \varepsilon \ , \tag{11}$$

where $A := A(d) \in \mathbb{R}^{N\times M}$ and $\varepsilon \sim \mathcal{N}(0, \Sigma_e)$. Assuming a normal distribution for the prior on ϑ,

$$p(\vartheta) = \mathcal{N}(\vartheta|\mu_\pi, \Sigma_\pi) \ , \tag{12}$$

we have seen in class that the posterior distribution is given by,

$$p(\vartheta|y, d) = \mathcal{N}(\vartheta|\mu_\vartheta, \Sigma_\vartheta) \ , \tag{13}$$

where the mean is given by,

$$\mu_\vartheta = \Sigma_\vartheta \Big(A^\top \Sigma_e^{-1} y + \Sigma_\pi^{-1}\mu_\pi \Big) \ , \tag{14}$$

and the covariance matrix by,

$$\Sigma_\vartheta = \Big(A^\top \Sigma_e^{-1} A + \Sigma_\pi^{-1} \Big)^{-1} \ . \tag{15}$$

Using (4), the information entropy of ϑ and $\vartheta|y, d$ is given by

$$H(\vartheta) = \frac{1}{2}\log\Big((2\pi \mathrm{e})^M |\Sigma_\pi|\Big) \ , \tag{16}$$

and

$$\begin{aligned} H(\vartheta|y, d) &= \frac{1}{2}\log\Big((2\pi \mathrm{e})^M |\Sigma_\vartheta|\Big) \\ &= \frac{1}{2}\log\Big((2\pi \mathrm{e})^M |A^\top \Sigma_e^{-1} A + \Sigma_\pi^{-1}|^{-1}\Big) \ . \end{aligned} \tag{17}$$

Notice that $H(\vartheta)$ does not depend on the design d and $H(\vartheta|y, d)$ does not depend on the measurements y. Substituting (16) and (17) in (10) we get,

$$\begin{aligned} U(d) &= -\int \frac{1}{2}\log\Big((2\pi \mathrm{e})^M |\Sigma_\vartheta|\Big)\, p(y|d)\, \mathrm{d}y - \frac{1}{2}\log\Big((2\pi \mathrm{e})^M |\Sigma_\pi|\Big) \\ &= \frac{1}{2}\log|A^\top \Sigma_e^{-1} A + \Sigma_\pi^{-1}| \, + \, c \ , \end{aligned} \tag{18}$$

where c is constant with respect to d. The utility function has an analytical expression which depends only the matrix, $A(d)$ defining the linear model, the covariance of the normal prior, Σ_π, and the prediction error covariance matrix, Σ_e. Thus, it is not required to perform simulations in order to evaluate the utility function. The optimal experimental design problem is thus considerably simplified.