6 Non-linear models

In physical modelling non-linearity is important: friction terms, for example, are inherently non-linear. Non-linearity also occurs when variables are transformed to modify their properties. Suppose that, rather than the linear form we are now familiar with, our model is

$$
\begin{aligned}
y_t &= f_t(\theta_t) + \nu_t & \nu_t &\sim \langle 0,\, V_t \rangle \\
\theta_t &= g_t(\theta_{t-1}) + \omega_t & \omega_t &\sim \langle \mathbf{0},\, W_t \rangle
\end{aligned}
$$

where θ_t is a n-vector, and $f_t : \mathbb{R}^n \to \mathbb{R}$ and $g_t : \mathbb{R}^n \to \mathbb{R}^n$ are known but not-necessarily-linear functions. We no longer have a simple rule for passing from the mean and variance of $\theta_{t-1} \mid D_{t-1}$ to the mean and variance of $\theta_t \mid D_{t-1}$ and from this to the mean and variance of $y_t \mid D_{t-1}$.

One situation in which non-linear g_t functions occur is when the state vector evolves according to an Ordinary Differential Equation (ODE):

$$
\frac{d\theta_i}{dt} = f_i(t, \theta) \qquad i = 1, \ldots, n
$$

for some known functions $f_1, \ldots, f_n$ (not the same f as in the observation equation). There are a number of methods by which we can solve these equations from a given $\theta(t)$ to find a good approximation for $\theta(t+1)$—the Runge-Kutta fourth order method is a very popular approach you may have come across—but they all amount to performing a non-linear operation involving evaluations of the f_i functions. So they can all be summarised in the form $\theta_t = g_t(\theta_{t-1})$, although we may not always have a simple form for g_t. Often we are unsure about θ_0, or we think that uncertainty will be introduced into the evolution of θ_t, or we want to account explicitly for the truncation errors in our solution method. For these reasons we would like to embed the ODE into a DLM to propogate uncertainty.

There are two approaches to handling non-linearity, depending on the problem. If f_t and g_t are simple differentiable functions and the non-linearity is not severe, then we can use linearisation. In other cases, where it is not possible to differentiate f_t or g_t or the non-linearity is considered to be substantial, we can use a sampling approach.

6.1 Linearisation

The linearisation approach is fairly standard, and often goes by the name 'generalised Kalman filter'. Consider the case of a non-linear but differentiable state equation. At time t we may compute the Taylor series expansion of $g_t(\theta_{t-1})$ around the known mean $m_{t-1} = \mathsf{E}[\theta_{t-1} \mid D_{t-1}]$ to get

$$
\theta_t = g_t(m_{t-1}) + \nabla g_t(m_{t-1}) \left(\theta_{t-1} - m_{t-1} \right) + \cdots + \omega_t
$$

where $\nabla g_t(m_{t-1})$ is the $n \times n$ gradient matrix of $g_t(\theta)$ evaluated at $\theta = m_{t-1}$. That is, $\left[\nabla g_t(m_{t-1}) \right]_{ij}$ is the derivative of θ_{ti} w.r.t. $\theta_{t-1,j}$, evaluated at $\theta_{t-1} = m_{t-1}$. Truncating after the linear term and writing $G_t = \nabla g_t(m_{t-1})$, we have the approximately linear model

$$
\theta_t \approx h_t + G_t\, \theta_{t-1} + \omega_t
$$

where $h_t = g_t(m_{t-1}) - G_t\, m_{t-1}$. The truncation will only be appropriate if the non-linear terms in the expansion of g_t die away rapidly. In this case, we have

$$
\begin{aligned}
a_t &= \mathsf{E}[\theta_t \mid D_{t-1}] \approx g_t(m_{t-1}) \\
R_t &= \mathsf{Var}\left[\theta_t \mid D_{t-1}\right] \approx G_t\, C_{t-1}\, (G_t)^\intercal + W_t
\end{aligned}
$$

like before, although we may want to boost W_t in order to account for the truncation error.

Exactly the same approach can be used for the observation equation, although now we expand $f_t(\theta_t)$ around $a_t = \mathsf{E}[\theta_t \mid D_{t-1}]$:

$$y_t = f_t(a_t) + \nabla f_t(a_t)(\theta_t - a_t) + \cdots + \nu_t$$

($\nabla f_t(a_t)$ is a row-vector) and we set $(F_t)^\intercal = \nabla f_t(a_t)$, and find that $\mathsf{E}[y_t \mid D_{t-1}] \approx f_t(a_t)$ and $\mathsf{Var}\,[y_t \mid D_{t-1}] \approx (F_t)^\intercal R_t\, F_t + V_t$ where, once again, we might want to bump V_t up a bit.

6.2 Sampling

What happens when f_t and/or g_t are hard to differentiate, or the non-linearity is thought to be substantial? This is a developing research area with several strands, but the following approach is promising. As it is particularly applicable to ODEs I will describe it for the state equation, but the same steps are followed for the observation equation.

The key idea is that rather than find the mean and variance of $\theta_t \mid D_{t-1}$ by looking at g_t only at $\theta_{t-1} = m_{t-1}$, we look more widely in the region around m_{t-1}. Naturally this is more work as we have to evaluate g_t more often, but it does not require an explicit calculation of the gradient matrix, and you might think that by looking in more detail at g_t we are better able to understand how it maps the mean and variance of θ_{t-1}.

Before we go on, a bit of matrix revision. Consider the relation $Z = XY^\intercal$. We know that $Z_{ij} = \sum_k X_{ik}\, Y_{jk}$. This shows that there is another way to write Z, in terms of the columns of X and Y:

$$Z = \sum_k X_{:k}(Y_{:k})^\intercal$$

where in general $A_{:k}$ is column k of matrix A. Now let's consider the spectral decomposition of $C_{t-1} = \mathsf{Var}\,[\theta_{t-1} \mid D_{t-1}]$, which we write as $C_{t-1} = \Gamma\Lambda\Gamma^\intercal$. By taking $Q = \Gamma\Lambda^{1/2}$ we find that

$$C_{t-1} = QQ^\intercal = \sum_k Q_{:k}(Q_{:k})^\intercal,$$

which is the key to the following approach.

We are going to evaluate g_t in total $2n$ times. We collect the evaluation points as the rows of the matrix X. Denoting row i (as a column vector) as $X_{i:}$, our $2n$ rows are

$$X_{i:} = (-1)^i \times \sqrt{n} \times Q_{:j} \qquad i = 1, \dots, 2n, \quad j := 1 + \lfloor (i-1)/2 \rfloor$$

where $\lfloor x \rfloor$ is the largest integer less than x. Note that rows i and $i+1$ for odd i both use *column* j of Q (rows 1 and 2 use column 1, rows 3 and 4 use column 2, and so on). Obviously the mean value of the rows of X is zero, since each pair of rows cancel out. But perhaps less obviously the variance of the rows is equal to C_{t-1}:

$$\frac{1}{2n} \sum_{1=1}^{2n} X_{i:}\,(X_{i:})^\intercal = \frac{n}{2n} \sum_{1=1}^{2n} Q_{:j}\,(Q_{:j})^\intercal = \frac{1}{2}(2\,C_{t-1}) = C_{t-1}.$$

So if we were to add m_{t-1} to each of the rows of X then we would have a design of $2n$ points with mean m_{t-1} and variance C_{t-1}.

To approximate $a_t = \mathsf{E}[\theta_t \mid D_{t-1}]$ and $R_t = \mathsf{Var}\,[\theta_t \mid D_{t-1}]$ we evaluate the function g_t at each of the rows of X. Call the result Y, with

$$Y_{i:} = g_t(m_{t-1} + X_{i:}) \qquad i = 1, \dots, 2n.$$

To find the mean a_t we compute the sample mean of the output data

$$a_t = \mathsf{E}[\theta_t \mid D_{t-1}] \approx \frac{1}{2n} \sum_{i=1}^{2n} Y_{i:}$$

while to find the variance R_t we compute the variance of the output data, and then add on $W_t = \mathsf{Var}\,[\omega_t]$:

$$R_t = \mathsf{Var}\,[\theta_t \mid D_{t-1}] \approx \frac{1}{2n} \sum_{i=1}^{2n} (Y_{i:} - a_t)(Y_{i:} - a_t)^\intercal + W_t.$$

We can check that this approach will give identical results to the standard approach in the case where g_t is linear. For then $Y_{i:} = G_t\, X_{i:}$ and it follows immediately that

$$a_t = \frac{1}{2n} \sum_{i=1}^{2n} \left(G_t\, X_{i:} \right) = G_t \left(\frac{1}{2n} \sum_{i=1}^{2n} X_{i:} \right) = G_t\, m_{t-1},$$

and there is a similar result for the variance. But the great strength of this approach is that we can use it for any g_t, including what is known as a 'black box': a function about which we know nothing at all.

More details can be found in S. Julier, J. Uhlmann and H. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators", *IEEE Transactions on Automatic Control*, 2000, **45**, pp. 477-82.

Exercises

1. Consider the case of an object in 2-dimensional Euclidean space falling freely (but not necessarily vertically) under gravity. Denote the location of this object by (x_t, y_t), and denote the general state vector as θ_t.

   (a) What is an appropriate state equation for this object? How do you interpret the elements of the state vector?

   (b) What factors might you consider in determining the error variance for your state equation?

   (c) Every five seconds I record the angle α_t between the ground and the object from my location $(\bar{x}, 0)$. What is the observation equation?

   (d) Given $\theta_t \mid D_{t-1} \sim \langle a_t,\, R_t \rangle$, how do I compute the mean and variance of $y_t \mid D_{t-1}$, and the covariance $\mathsf{Cov}[\theta_t,\, y_t \mid D_{t-1}]$?