BAYESIAN TOMOGRAPHY

John Skilling (john@skilling.co.uk)

Paris, March 2015

Abstract

Tomography measures the density of a body (usually by its opacity to X-rays) along multiple lines of sight.

$$\text{Body} \xrightarrow{\text{X-rays}} \text{Opacity data}$$

From these line integrals, we reconstruct an image $\rho(\mathbf{x})$ of the density as a function of position. The user then interprets this in terms of material classification (such as bone or muscle or fat or fluid or cavity).

$$\text{Opacity data} \xrightarrow{\text{invert}} \text{Density image} \xrightarrow{\text{interpret}} \text{Material model}$$

Data are naturally incomplete and noisy, which implies that our inferences must be similarly imprecise. Data do not give us a single image $\widehat{\rho}$. Instead, data give us a *probability distribution* $\text{Prob}(\rho)$ of plausible images. This uniquely correct approach is called *Bayesian* (even though that is historically unjust to Laplace).

Bayesian analysis starts with the prior probability $\text{Prior}(\rho)$ of those images that might originally seem plausible. That prior is modulated by the data (through the *likelihood function*) to give the resulting distribution $\text{Posterior}(\rho)$ of those images that remain plausible *after the data are used*.

$$\text{Posterior} \propto \text{Prior} \times \text{Likelihood}$$

Usually, only a tiny proportion $O(e^{-\text{ size of dataset}})$ of prior possibilities survives into the posterior, so that Bayesian analysis was essentially impossible before computers. Even with computers, direct search is impossible in large dimension, and we need numerical methods such as nested sampling to guide the exploration. But this can now be done, easily and generally. The exponential curse of dimensionality is removed.

Bayesian analysis can alternatively start directly with a prior over plausible material models. It does not need to pass through the intermediate step of density.

$$\text{Opacity data} \xrightarrow{\qquad\qquad\text{Bayes}\qquad\qquad} \text{Material model}$$

This simplifies the approach and should make the results easier to use. Moreover, this trick enables *data fusion*, where different modalities (such as CT, MRI, ultrasound) can all be brought to bear on the *same* material model, which will become clearer and more informative than any individual reconstruction.

The Bayesian approach thus suggests a fundamental reformulation of imagery in general, and tomography in particular.