An approach to sensor fusion using geostatistics

Ben Ingram[1], Dan Cornford[2]

[1]Aston University
B.R.Ingram@aston.ac.uk

[2] Aston University
d.cornford@aston.ac.uk

Large networks of sensors are becoming increasingly common due to advances in inexpensive sensor technology. These large networks are often comprised of many smaller sensor networks. The specifics of such sub-networks can vary depending on a number of factors such as: the year it was deployed, politics, cost, geographical location, etc... As a result, in real world network monitoring situations there tends to be a large number of diverse sensor types, and hence, observations are heterogeneous, both in terms of the detail of their observation characteristics, and their error characteristics. The presence of heterogeneities in the characteristics of different sensors can present a number of problems when trying to determine covariance, or variogram, parameters in a geostatistical analysis. This is because different sensor models will produce different observations of the same underlying process, which means sample variograms cannot be constructed. A local inversion of the sensor model will not be optimal, and could produce complex error distributions on the retrieved process. The natural solution is a model based approach which leads naturally to a Bayesian formulation of the interpolation problem, with a Gaussian process prior placed over the latent, underlying process. This allows decision makers use such diverse observations to generate coherent maps of an area of interest to make informed decisions about potentially life threatening scenarios, taking into account the dominant sources of uncertainty.

In this poster we consider a model based geostatistical approach to sensor fusion. A coherent model of the observations is constructed using the characteristic response of each sensor to the underlying, latent process, in terms of an error distribution and some (potentially non-linear) observation operator. The inclusion of the sensor response in the model is enabled by the specification of arbitrary, and potentially different, likelihood models for each observation. We show that in the presence of observations collected from unreliable sources, robust likelihood models can be used to increase the stability of the model. Finally, when the number of observations becomes prohibitively large (more than a few thousand), we show how the computational complexity of the model can be reduced without significant loss of prediction accuracy.