The role of mechanistic models in Bayesian inference

Dan Cornford[1], Alexis Boukouvalas, Yuan Shen, Michael Vrettas, Manfred Opper and many others

Neural Computing Research Group

Aston University
Engineering & Applied Science

Grasmere, Sept 2008, BARK [2]

[1] http://wiki.aston.ac.uk/DanCornford/
[2] WARNING: This talk may contain traces of machine learning.

I am interested in problems in the 'real world'

- weather forecasting - will it rain tomorrow?
- climate - should I really move to higher ground?
- disease modelling - will rabies become established in Finland?
- environmental monitoring - is it safe to eat those mushrooms?

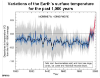

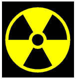

- These are in essence difficult regression (decision) problems – where prior knowledge is key.
- Prior knowledge is often in the form of physical laws, typically implemented as simulators.

Here is my naive view of the problem.

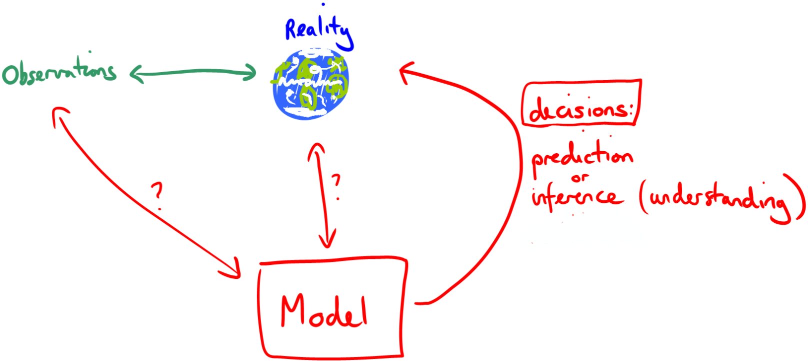

Simulators and modelling

- In the physical and natural sciences the emphasis has been on laws, processes and mechanisms.
- These are combined in complex models, which I will call simulators, almost always implemented as computer code.
- Historically these models were deterministic – a given input produces a given output.
- Climate models are great examples – vastly complex, running on the largest computers in the world, and barely using data!
- So how do these (deterministic) models fit in a Bayesian framework?

Simulators and modelling

- In the physical and natural sciences the emphasis has been on laws, processes and mechanisms.
- These are combined in complex models, which I will call simulators, almost always implemented as computer code.
- Historically these models were deterministic – a given input produces a given output.
- Climate models are great examples – vastly complex, running on the largest computers in the world, and barely using data!
- So how do these (deterministic) models fit in a Bayesian framework?
- That's the topic of this talk. (And the next N years of my life!)

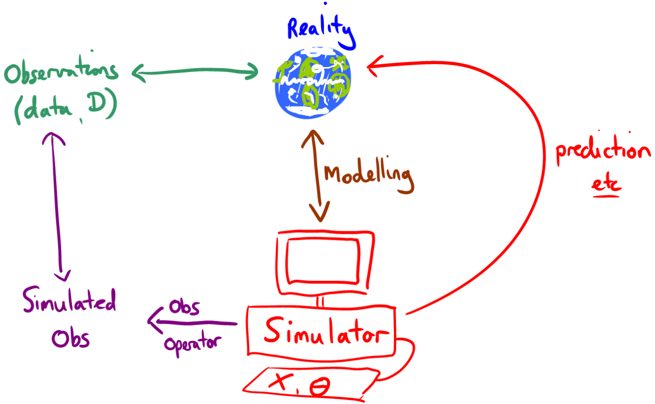

Aston University
Engineering & Applied Science

How can we use deterministic simulators?

Throwing them out is wasteful

- These models encode prior knowledge from many sources – we need to use this – especially in extrapolation situations.

Uncertainty analysis, sensitivity analysis:

- What is the impact of uncertainties in inputs on outputs?
- Assuming we can define distributions over the relevant inputs (elicitation), sample from these and use Monte Carlo methods.
- Analysis of the simulator - no use of data here!

- One problem is that (interesting) simulators are very expensive to run.
- A solution is to emulate the computer code, using e.g. a Gaussian process.

Aston University
Engineering & Applied Science

MUCM: Managing Uncertainty in Complex Models

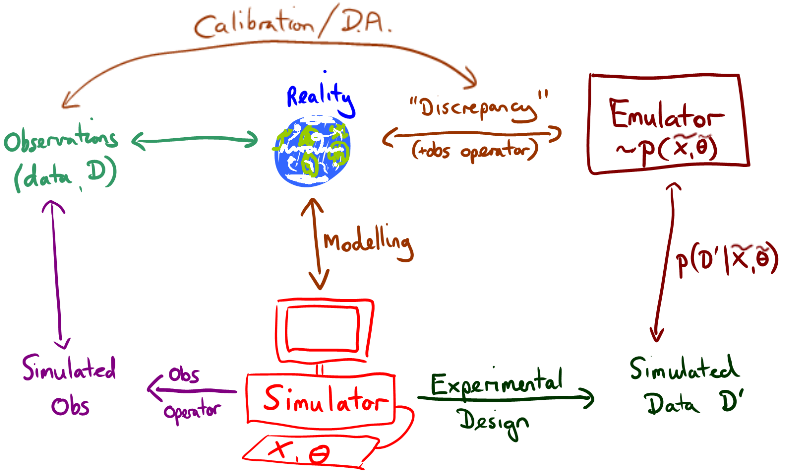

Jonty Rougier → Reification

MUCM: Managing Uncertainty in Complex Models

RCUK Basic Technology Project:

- Extend emulation to: high dimensional models, dynamic simulators, calibration, linking models to reality, improved designs.
- Make these methods accessible to applications.

This area has many open challenges; we are investigating:

- emulation of stochastic models, often with non-Gaussian output distributions.
- emulation of high dimensional models, using dimension reduction.

Aston University
Engineering & Applied Science

Stochastic simulators

To make further progress:

- Simulators must become stochastic – randomness arises from incomplete knowledge, not internally.
- Model error is critical to represent; simulators then define (complex) priors.
- The problem now is how to do inference with such complex priors.
- Full Monte Carlo (or MCMC) is not conceivable – these process based models are often very high dimensional, being based on partial / ordinary differential equations.
- So can we approximate somehow?
- Emulation remains an option.

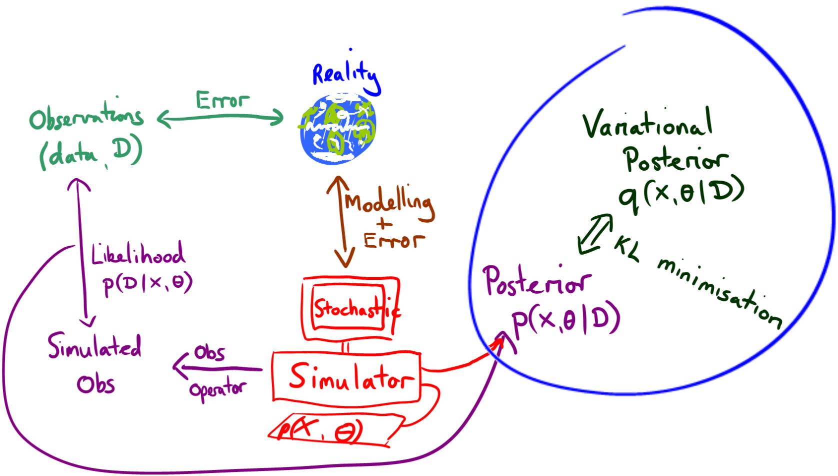
Reality
Observations
(data, D)
Error
Modelling
+
Error
Likelihood
p(D|x,θ)
Simulated
Obs
Obs
Operator
Stochastic
Simulator
p(x,θ)
Variational
Posterior
q(x,θ|D)
KL minimisation
Posterior
p(x,θ|D)

Summary

- Machine learning has shown that really difficult problems can be tackled and solved.
- A range of novel approximate inference methods, particularly variational methods have been developed.
- These could be applied to process driven models if we characterise the model errors, and emulation is an option too, with appropriate relation of the model to reality.
- The overall Bayesian framework stays the same, but the prior is more complex.
- Lots of unresolved issues with implementation – size / speed, optimisation, exploiting structure, and more.
- But if we can do it, we can use data **and** extrapolate well.