Bayesian Linear Uncertainty Analysis for complex computer models

Michael Goldstein
Department of Mathematical Sciences, Durham University *

*Thanks to Leverhulme for funding through the Durham Tipping points project, to EPSRC for funding on the Managing Uncertainty for Complex Models Consortium, to NERC for funding on the RAPID and PURE programmes

Complex physical models.

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators (like climate models).

Complex physical models.

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators (like climate models).

To use complex simulators to make statements about physical systems (like climate), we need to quantify the uncertainty involved in moving from the model to the system.

Complex physical models.

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators (like climate models).

To use complex simulators to make statements about physical systems (like climate), we need to quantify the uncertainty involved in moving from the model to the system.

This talk will give an overview of some important features of this area, from a Bayes linear viewpoint.

Complex physical models.

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators (like climate models).

To use complex simulators to make statements about physical systems (like climate), we need to quantify the uncertainty involved in moving from the model to the system.

This talk will give an overview of some important features of this area, from a Bayes linear viewpoint.

Camila’s talk (following) will apply these ideas, for a climate model which is simple but has interesting behaviour.

Examples

Oil reservoir Simulator inputs are properties of the reservoir (permeabilities, porosities, faults). Outputs are behaviour at wells (gas/oil production, water cut). The aim is to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Examples

Oil reservoir Simulator inputs are properties of the reservoir (permeabilities, porosities, faults). Outputs are behaviour at wells (gas/oil production, water cut). The aim is to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Galaxy formation The study of the development of the Universe is carried out by using a Galaxy formation simulator. The inputs are certain underlying physical constants (considered true but unknown). The outputs are the configuration of galaxies in the universe. The aim is to gain information about the physical processes underlying the Universe.

Examples

Oil reservoir Simulator inputs are properties of the reservoir (permeabilities, porosities, faults). Outputs are behaviour at wells (gas/oil production, water cut). The aim is to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Galaxy formation The study of the development of the Universe is carried out by using a Galaxy formation simulator. The inputs are certain underlying physical constants (considered true but unknown). The outputs are the configuration of galaxies in the universe. The aim is to gain information about the physical processes underlying the Universe.

Climate change Global climate simulators are used to assess likely effects of human intervention upon future climate behaviour. Inputs are physical constants describing the evolution of climate in response to properties like CO2 forcing. Outputs are features of global future climate. Aims are scientific - to learn about large scale interactions which determine climate - and practical, as such simulators provide evidence for the need to change our behaviour before irreversible changes are set into motion.

Sources of Uncertainty.

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

(iv) stochastic uncertainty (either the model is stochastic, or it should be),

(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

(vi) structural uncertainty (the model only approximates the physical system),

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

(iv) stochastic uncertainty (either the model is stochastic, or it should be),

(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

(vi) structural uncertainty (the model only approximates the physical system),

(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

(iv) stochastic uncertainty (either the model is stochastic, or it should be),

(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

(vi) structural uncertainty (the model only approximates the physical system),

(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),

(viii) multi-model uncertainty (usually we have not one but many models related to the physical system)

Sources of Uncertainty.

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).
(vi) structural uncertainty (the model only approximates the physical system),
(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),
(viii) multi-model uncertainty (usually we have not one but many models related to the physical system)
(ix) decision uncertainty (to use the model to influence real world outcomes, we need to relate things in the world that we can influence to inputs to the simulator and through outputs to actual impacts. These links are uncertain.)

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the "appropriate" (in some sense) choice, x^*, for the system properties x,

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the "appropriate" (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the “appropriate” (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

the use that we can make of historical observations z, observed with error on a subset y_h of y, both to test and to constrain the model,

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the “appropriate” (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

the use that we can make of historical observations z, observed with error on a subset y_h of y, both to test and to constrain the model,

the optimal assignment of any decision inputs, d, in the model.

General form

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where

x: input vector, representing unknown properties of the physical system;

$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the “appropriate” (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

the use that we can make of historical observations z, observed with error on a subset y_h of y, both to test and to constrain the model,

the optimal assignment of any decision inputs, d, in the model.

[In a climate model, y_h might correspond to historical climate outcomes over space and time, y to current and future climate, and the “decisions” might correspond to different policy relevant choices such as carbon emission scenarios.]

“Solving” these problems

If observations, z, are made without error and the model is perfect reproduction of the system, we can write $z = f_h(x^*)$, invert f_h to find x^*, learn about all future components of $y = f(x^*)$ and choose decision elements of x^* to optimise properties of y.

“Solving” these problems

If observations, z, are made without error and the model is perfect reproduction of the system, we can write $z = f_h(x^*)$, invert f_h to find x^*, learn about all future components of $y = f(x^*)$ and choose decision elements of x^* to optimise properties of y.
COMMENT: This would be very hard.

“Solving” these problems

If observations, z, are made without error and the model is perfect reproduction of the system, we can write $z = f_h(x^*)$, invert f_h to find x^*, learn about all future components of $y = f(x^*)$ and choose decision elements of x^* to optimise properties of y.
COMMENT: This would be very hard.
In practice, the observations z are made with error, and model is not the same as physical system so we must separate the uncertainty representation into two relations and carry out statistical inversion/optimisation:

$$z = y_h \oplus e, \;\; y = f(x^*) \oplus \epsilon$$

where e, ϵ have some appropriate probabilistic specification, possibly involving parameters which require estimation.
($a \oplus b$ means $a + b$ where a, b are independent.)

“Solving” these problems

If observations, z, are made without error and the model is perfect reproduction of the system, we can write $z = f_h(x^*)$, invert f_h to find x^*, learn about all future components of $y = f(x^*)$ and choose decision elements of x^* to optimise properties of y.

COMMENT: This would be very hard.

In practice, the observations z are made with error, and model is not the same as physical system so we must separate the uncertainty representation into two relations and carry out statistical inversion/optimisation:

$$z = y_h \oplus e, \;\; y = f(x^*) \oplus \epsilon$$

where e, ϵ have some appropriate probabilistic specification, possibly involving parameters which require estimation.

($a \oplus b$ means $a + b$ where a, b are independent.)

COMMENT: This is much harder.

“Solving” these problems

If observations, z, are made without error and the model is perfect reproduction of the system, we can write $z = f_h(x^*)$, invert f_h to find x^*, learn about all future components of $y = f(x^*)$ and choose decision elements of x^* to optimise properties of y.
COMMENT: This would be very hard.
In practice, the observations z are made with error, and model is not the same as physical system so we must separate the uncertainty representation into two relations and carry out statistical inversion/optimisation:

$$z = y_h \oplus e, \;\; y = f(x^*) \oplus \epsilon$$

where e, ϵ have some appropriate probabilistic specification, possibly involving parameters which require estimation.
($a \oplus b$ means $a + b$ where a, b are independent.)
COMMENT: This is much harder.
COMMENT And we still haven’t accounted for condition uncertainty, multi-model uncertainty, etc.

RAPID-WATCH

What are the implications of RAPID-WATCH observing system data and other recent observations for estimates of the risk due to rapid change in the MOC? In this context risk is taken to mean the probability of rapid change in the MOC and the consequent impact on climate (affecting temperatures, precipitation, sea level, for example). This project must:

* contribute to the MOC observing system assessment in 2011;
* investigate how observations of the MOC can be used to constrain estimates of the probability of rapid MOC change, including magnitude and rate of change;
* make sound statistical inferences about the real climate system from model simulations and observations;
* investigate the dependence of model uncertainty on such factors as changes of resolution;
* assess model uncertainty in climate impacts and characterise impacts that have received less attention (eg frequency of extremes).

The project must also demonstrate close partnership with the Hadley Centre.

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

In this interpretation, any probability statement is the judgement of a named individual, so we should speak not of “the probability of rapid climate change”, but instead of this expert’s probability or that expert’s probability (or this group of experts shared probability) of rapid climate change.

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

In this interpretation, any probability statement is the judgement of a named individual, so we should speak not of “the probability of rapid climate change”, but instead of this expert’s probability or that expert’s probability (or this group of experts shared probability) of rapid climate change.

So, a natural objective of a scientifically rigorous uncertainty analysis should be probabilities which are

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

In this interpretation, any probability statement is the judgement of a named individual, so we should speak not of “the probability of rapid climate change”, but instead of this expert’s probability or that expert’s probability (or this group of experts shared probability) of rapid climate change.

So, a natural objective of a scientifically rigorous uncertainty analysis should be probabilities which are

asserted by at least one person, expert in the area

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

In this interpretation, any probability statement is the judgement of a named individual, so we should speak not of “the probability of rapid climate change”, but instead of this expert’s probability or that expert’s probability (or this group of experts shared probability) of rapid climate change.

So, a natural objective of a scientifically rigorous uncertainty analysis should be probabilities which are

asserted by at least one person, expert in the area

for reasons that are open to outside scrutiny

Subjectivist Bayes

In the subjectivist Bayes view, a probability statement is the uncertainty judgement of a specified individual, expressed on the scale of probability. This interpretation has a well-defined and operational meaning.

In this interpretation, any probability statement is the judgement of a named individual, so we should speak not of “the probability of rapid climate change”, but instead of this expert’s probability or that expert’s probability (or this group of experts shared probability) of rapid climate change.

So, a natural objective of a scientifically rigorous uncertainty analysis should be probabilities which are

asserted by at least one person, expert in the area

for reasons that are open to outside scrutiny

including consideration of the range of alternative probability judgements which could plausibly be reached by other experts in the area

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.
We make expectation, rather than probability, the primitive for the theory, following de Finetti “Theory of Probability”(1974,1975).

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.

We make expectation, rather than probability, the primitive for the theory, following de Finetti “Theory of Probability”(1974,1975).

He chooses expectation over probability as, if expectation is primitive, then we can choose to make as many or as few expectation statements as we choose, whereas, if probability is primitive, then we must make all of the probability statements before we can make any of the expectation statements.

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.

We make expectation, rather than probability, the primitive for the theory, following de Finetti “Theory of Probability”(1974,1975).

He chooses expectation over probability as, if expectation is primitive, then we can choose to make as many or as few expectation statements as we choose, whereas, if probability is primitive, then we must make all of the probability statements before we can make any of the expectation statements.

(As any probability is the expectation of the indicator function for the corresponding event, we make all of the probability statements that we need directly as expectations.)

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.

We make expectation, rather than probability, the primitive for the theory, following de Finetti "Theory of Probability"(1974,1975).

He chooses expectation over probability as, if expectation is primitive, then we can choose to make as many or as few expectation statements as we choose, whereas, if probability is primitive, then we must make all of the probability statements before we can make any of the expectation statements.

(As any probability is the expectation of the indicator function for the corresponding event, we make all of the probability statements that we need directly as expectations.)

Analysis based on expectation as primitive is termed (by me) Bayes linear analysis (because of the linearity properties of expectation).

Expectation as primitive

It is hard to carry out a full subjectivist Bayes analysis for a complex computer model, because of the extremely finely detailed level at which our judgements must be converted to uncertainties. However, we do have a choice.

We make expectation, rather than probability, the primitive for the theory, following de Finetti "Theory of Probability"(1974,1975).

He chooses expectation over probability as, if expectation is primitive, then we can choose to make as many or as few expectation statements as we choose, whereas, if probability is primitive, then we must make all of the probability statements before we can make any of the expectation statements.

(As any probability is the expectation of the indicator function for the corresponding event, we make all of the probability statements that we need directly as expectations.)

Analysis based on expectation as primitive is termed (by me) Bayes linear analysis (because of the linearity properties of expectation).

For a detailed account of this approach see

Bayes linear Statistics: Theory and Methods, 2007, Wiley

Michael Goldstein and David Wooff

Adjusted expectation

We can analyse expectations directly using the Bayes linear approach, in which we make direct prior specifications for that collection of means, variances and covariances which we are able to assess, and update these by linear fitting

Adjusted expectation

We can analyse expectations directly using the Bayes linear approach, in which we make direct prior specifications for that collection of means, variances and covariances which we are able to assess, and update these by linear fitting

Suppose that we have two collections of random quantities, namely vectors $\boldsymbol{B} = (B_1, ..., B_r)$, $\boldsymbol{D} = (D_0, D_1, ..., D_s)$ where $D_0 = 1$. We observe $\boldsymbol{D}$

Adjusted expectation

We can analyse expectations directly using the Bayes linear approach, in which we make direct prior specifications for that collection of means, variances and covariances which we are able to assess, and update these by linear fitting

Suppose that we have two collections of random quantities, namely vectors $\boldsymbol{B} = (B_1, ..., B_r)$, $\boldsymbol{D} = (D_0, D_1, ..., D_s)$ where $D_0 = 1$. We observe $\boldsymbol{D}$

The *adjusted* or *Bayes linear* expectation and variance for $\boldsymbol{B}$ given $\boldsymbol{D}$ are

$$\mathsf{E}_D[B] = \mathrm{E}(B) + \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}(D - \mathrm{E}(D)),$$
$$\mathsf{Var}_D[B] = \mathrm{Var}(B) - \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}\mathrm{Cov}(D, Z)$$

Adjusted expectation

We can analyse expectations directly using the Bayes linear approach, in which we make direct prior specifications for that collection of means, variances and covariances which we are able to assess, and update these by linear fitting

Suppose that we have two collections of random quantities, namely vectors $\boldsymbol{B} = (B_1, ..., B_r)$, $\boldsymbol{D} = (D_0, D_1, ..., D_s)$ where $D_0 = 1$. We observe $\boldsymbol{D}$

The *adjusted* or *Bayes linear* expectation and variance for $\boldsymbol{B}$ given $\boldsymbol{D}$ are

$$\mathsf{E}_D[B] = \mathrm{E}(B) + \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}(D - \mathrm{E}(D)),$$
$$\mathsf{Var}_D[B] = \mathrm{Var}(B) - \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}\mathrm{Cov}(D, Z)$$

Working with expectations has the advantages that the uncertainty specification is simpler, the analysis is much faster and more straightforward and there are rigorous foundations for the approach.

Adjusted expectation

We can analyse expectations directly using the Bayes linear approach, in which we make direct prior specifications for that collection of means, variances and covariances which we are able to assess, and update these by linear fitting

Suppose that we have two collections of random quantities, namely vectors $\boldsymbol{B} = (B_1, ..., B_r)$, $\boldsymbol{D} = (D_0, D_1, ..., D_s)$ where $D_0 = 1$. We observe $\boldsymbol{D}$

The *adjusted* or *Bayes linear* expectation and variance for $\boldsymbol{B}$ given $\boldsymbol{D}$ are

$$\mathsf{E}_D[B] = \mathrm{E}(B) + \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}(D - \mathrm{E}(D)),$$
$$\mathsf{Var}_D[B] = \mathrm{Var}(B) - \mathrm{Cov}(B, D)\mathrm{Var}(D)^{-1}\mathrm{Cov}(D, Z)$$

Working with expectations has the advantages that the uncertainty specification is simpler, the analysis is much faster and more straightforward and there are rigorous foundations for the approach.

In practice, we may choose a mix of expectation and probability based methods. Here our focus is on the expectation based methodology.

Limitations of physical models

A physical model describes how
system properties (the inputs to the model)
affect system behaviour (the output of the model)
using two types of simplification:

Limitations of physical models

A physical model describes how
system properties (the inputs to the model)
affect system behaviour (the output of the model)
using two types of simplification:

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

Limitations of physical models

A physical model describes how
system properties (the inputs to the model)
affect system behaviour (the output of the model)
using two types of simplification:

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

(ii) we approximate the rules for finding system behaviour given system properties (because of necessary mathematical simplifications, simplifications for numerical tractability, and because we do not fully understand the physical laws which govern the process).

Limitations of physical models

A physical model describes how
system properties (the inputs to the model)
affect system behaviour (the output of the model)
using two types of simplification:

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

(ii) we approximate the rules for finding system behaviour given system properties (because of necessary mathematical simplifications, simplifications for numerical tractability, and because we do not fully understand the physical laws which govern the process).

Problems arise when we forget these simplifications and confuse the analysis of the model with the corresponding analysis for the physical system itself.

Structural uncertainty analysis

Models do not produce statements about reality, however carefully they are analysed. Such statements require **structural uncertainty** assessment, taking account the mismatch between the simulator and the physical system.

Structural uncertainty analysis

Models do not produce statements about reality, however carefully they are analysed. Such statements require **structural uncertainty** assessment, taking account the mismatch between the simulator and the physical system.

One of the simplest, and most popular, approaches is to suppose that there is an appropriate choice of system properties x^* (currently unknown), so that $f(x^*)$ contains all the information about the system:

$$y = f(x^*) \oplus \epsilon$$

where ϵ, the model or structural discrepancy, has some appropriate probabilistic specification, possibly involving parameters which require estimation, and is taken to be independent of f, x_0, e.

Structural uncertainty analysis

Models do not produce statements about reality, however carefully they are analysed. Such statements require **structural uncertainty** assessment, taking account the mismatch between the simulator and the physical system.

One of the simplest, and most popular, approaches is to suppose that there is an appropriate choice of system properties x^* (currently unknown), so that $f(x^*)$ contains all the information about the system:

$$y = f(x^*) \oplus \epsilon$$

where ϵ, the model or structural discrepancy, has some appropriate probabilistic specification, possibly involving parameters which require estimation, and is taken to be independent of f, x_0, e.
Careful structural uncertainty assessessment is crucial. Two aspects:
(i) Internal discrepancy: aspects we assess by simulator experiments
(ii) External discrepancy: inherent limitations of modelling process

Reification

What does the simulator f tell us about the system y?

Reification

What does the simulator f tell us about the system y?

Consider both our inputs x and the simulator f as abstractions/simplifications of real physical quantities and processes (through approximations in physics, solution methods, level of detail, limitations of current understanding).

Reification

What does the simulator f tell us about the system y?

Consider both our inputs x and the simulator f as abstractions/simplifications of real physical quantities and processes (through approximations in physics, solution methods, level of detail, limitations of current understanding).

These form approximations to a much more realistic simulator f^*, for which real, physical x^* would be the best input,
[in the sense that $(y - f^*(x^*))$ would be judged independent of (x^*, f^*).]

Reification

What does the simulator f tell us about the system y?

Consider both our inputs x and the simulator f as abstractions/simplifications of real physical quantities and processes (through approximations in physics, solution methods, level of detail, limitations of current understanding).

These form approximations to a much more realistic simulator f^*, for which real, physical x^* would be the best input,
[in the sense that $(y - f^*(x^*))$ would be judged independent of (x^*, f^*).]

We call f^* the reified simulator (from reify: to treat an abstract concept as if it was real).
Our model f is informative for y because f is informative for the more elaborate model f^* which is informative for y.

Reification

What does the simulator f tell us about the system y?

Consider both our inputs x and the simulator f as abstractions/simplifications of real physical quantities and processes (through approximations in physics, solution methods, level of detail, limitations of current understanding).

These form approximations to a much more realistic simulator f^*, for which real, physical x^* would be the best input,
[in the sense that $(y - f^*(x^*))$ would be judged independent of (x^*, f^*).]

We call f^* the reified simulator (from reify: to treat an abstract concept as if it was real).
Our model f is informative for y because f is informative for the more elaborate model f^* which is informative for y.
We can't evaluate f^*, but we can "emulate" it.

Relating the model and the system

Model evaluations

Actual system

System observations

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the evaluations to the notion of a 'best' evaluation
3. We link the 'best' evaluation to the actual system
4. We incorporate measurement error into the observations

Relating the model and the system

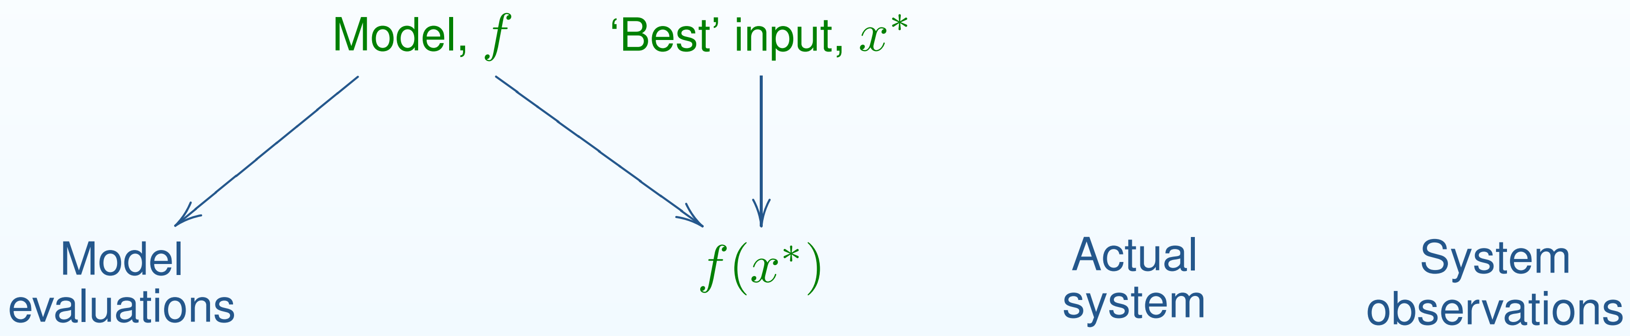

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the evaluations to the notion of a 'best' evaluation
3. We link the 'best' evaluation to the actual system
4. We incorporate measurement error into the observations

Relating the model and the system

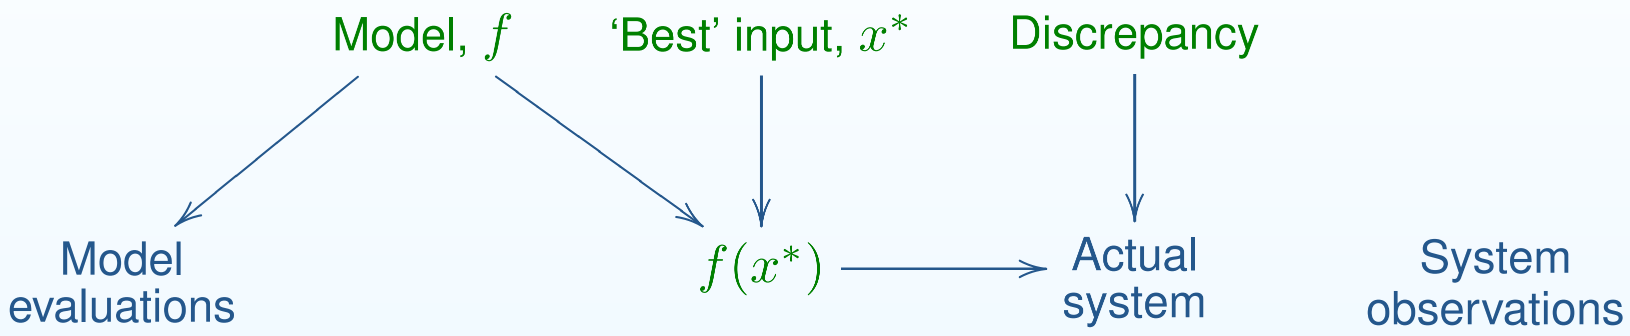

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the evaluations to the notion of a 'best' evaluation
3. We link the 'best' evaluation to the actual system
4. We incorporate measurement error into the observations

Relating the model and the system

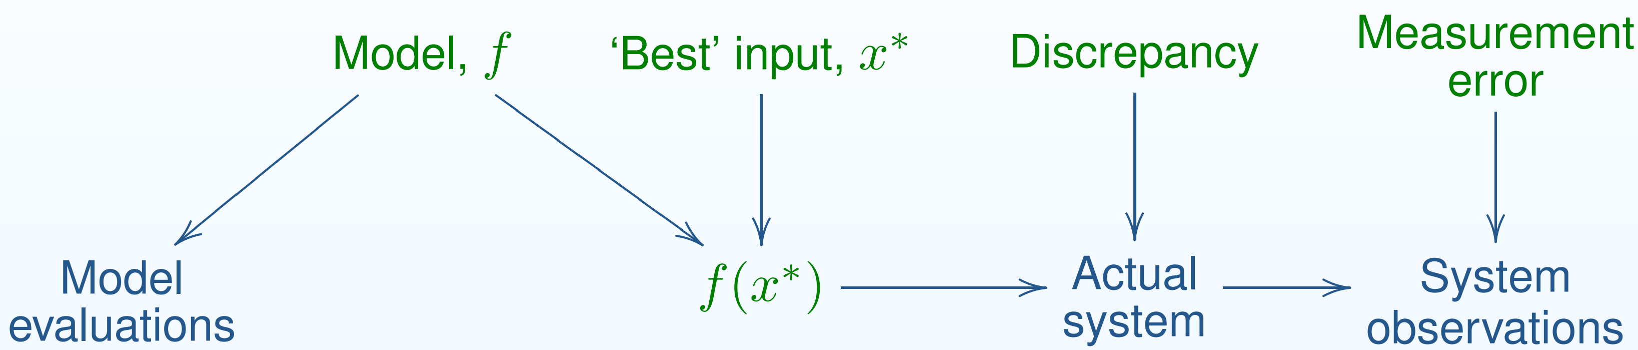

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the evaluations to the notion of a 'best' evaluation
3. We link the 'best' evaluation to the actual system
4. We incorporate measurement error into the observations

Relating the model and the system

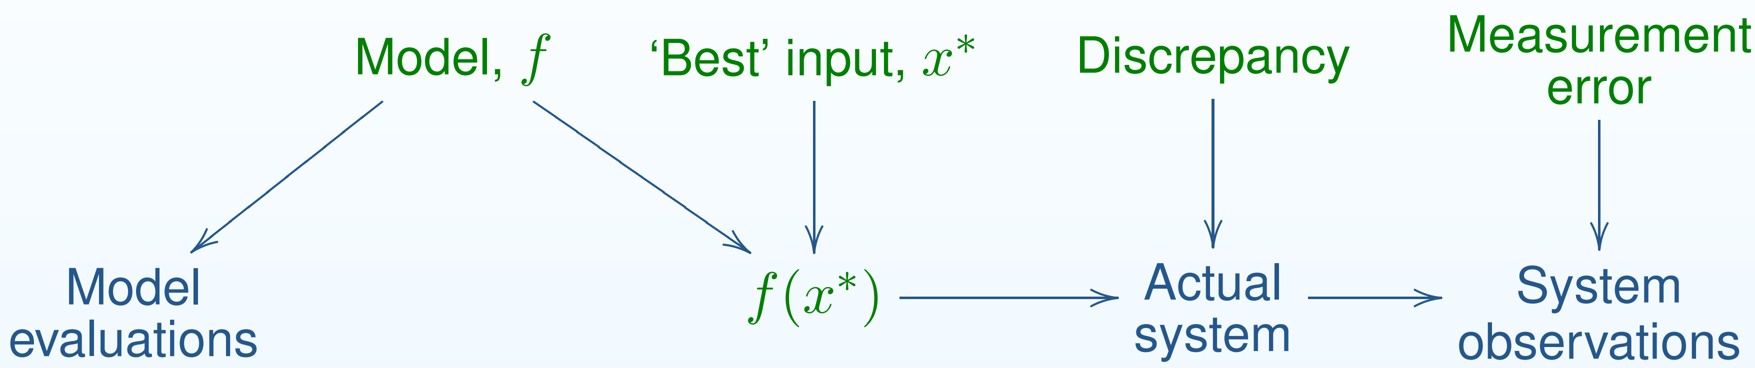

1. We start with the simple form for model discrepancy
2. We link the simulator to the reified form

Relating the model and the system

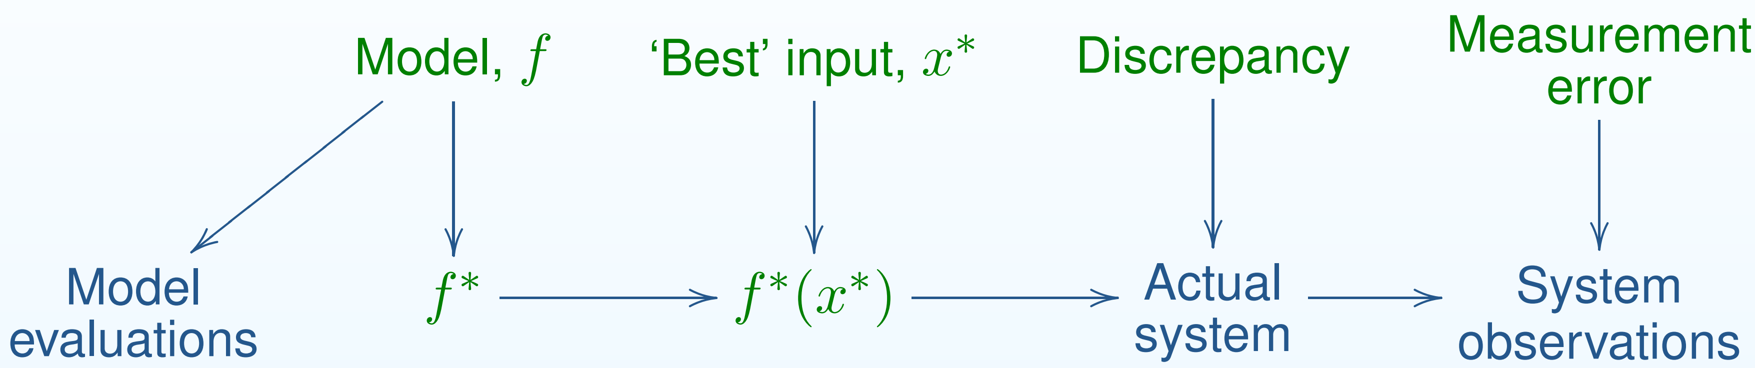

1. We start with the simple form for model discrepancy

2. We link the simulator to the reified form

Function emulation

Uncertainty analysis, for high dimensional problems, is even more challenging if the function $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x. [For example, large climate models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Function emulation

Uncertainty analysis, for high dimensional problems, is even more challenging if the function $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x. [For example, large climate models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.
Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.
Such a representation is often termed an emulator of the function - the emulator both suggests an approximation to the function and also contains an assessment of the likely magnitude of the error of the approximation.

Function emulation

Uncertainty analysis, for high dimensional problems, is even more challenging if the function $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x. [For example, large climate models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Such a representation is often termed an emulator of the function - the emulator both suggests an approximation to the function and also contains an assessment of the likely magnitude of the error of the approximation.

We use the emulator either to provide a full joint probabilistic description of all of the function values (full Bayes) or to assess expectations variances and covariances for pairs of function values (Bayes linear).

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

$$f_i(x) = \sum_j \beta_{ij} g_{ij}(x) \oplus u_i(x)$$

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

$$f_i(x) = \sum_j \beta_{ij} g_{ij}(x) \oplus u_i(x)$$

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, $u_i(x)$ is a weakly second order stationary stochastic process, with (for example) correlation function

$$\mathrm{Corr}(u_i(x), u_i(x')) = \exp(-(\tfrac{\|x-x'\|}{\theta_i})^2)$$

$Bg(x)$ expresses global variation in f. $u(x)$ expresses local variation in f

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

$$f_i(x) = \sum_j \beta_{ij} g_{ij}(x) \oplus u_i(x)$$

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, $u_i(x)$ is a weakly second order stationary stochastic process, with (for example) correlation function

$$\mathrm{Corr}(u_i(x), u_i(x')) = \exp(-(\tfrac{\|x-x'\|}{\theta_i})^2)$$

$Bg(x)$ expresses global variation in f. $u(x)$ expresses local variation in f
We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

$$f_i(x) = \sum_j \beta_{ij} g_{ij}(x) \oplus u_i(x)$$

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, $u_i(x)$ is a weakly second order stationary stochastic process, with (for example) correlation function

$$\mathrm{Corr}(u_i(x), u_i(x')) = \exp(-(\tfrac{\|x-x'\|}{\theta_i})^2)$$

$Bg(x)$ expresses global variation in f. $u(x)$ expresses local variation in f

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

We need careful (multi-output) experimental design to choose informative model evaluations, and detailed diagnostics to check emulator validity.

Linked emulators

If the simulator is really slow to evaluate, then we emulate by jointly modelling the simulator with a fast approximate version, f', plus older generations of the simulator which we've already emulated and so forth.

Linked emulators

If the simulator is really slow to evaluate, then we emulate by jointly modelling the simulator with a fast approximate version, f', plus older generations of the simulator which we've already emulated and so forth.

So, for example, based on many fast simulator evaluations, we build emulator

$$f_i'(x) = \sum_j \beta_{ij}' g_{ij}(x) \oplus u_i'(x)$$

Linked emulators

If the simulator is really slow to evaluate, then we emulate by jointly modelling the simulator with a fast approximate version, f', plus older generations of the simulator which we've already emulated and so forth.
So, for example, based on many fast simulator evaluations, we build emulator

$$f'_i(x) = \sum_j \beta'_{ij} g_{ij}(x) \oplus u'_i(x)$$

We use this form as the prior for the emulator for $f_i(x)$.
Then a relatively small number of evaluations of $f_i(x)$, using relations such as

$$\beta_{ij} = \alpha_i \beta'_{ij} + \gamma_{ij}$$

lets us adjust the prior emulator to an appropriate posterior emulator for $f_i(x)$.

History matching

Model calibration aims to identify “true” input parameters x^*. However

History matching

Model calibration aims to identify "true" input parameters x^*. However

(i) We may not believe in a unique true input value for the model;

History matching

Model calibration aims to identify "true" input parameters x^*. However

(i) We may not believe in a unique true input value for the model;

(ii) We may be unsure whether there are any good choices of input parameters (due to model deficiencies)

History matching

Model calibration aims to identify “true” input parameters x^*. However

(i) We may not believe in a unique true input value for the model;

(ii) We may be unsure whether there are any good choices of input parameters (due to model deficiencies)

(iii) Full Bayes calibration analysis may be very difficult/non-robust.

History matching

Model calibration aims to identify “true” input parameters x^*. However
(i) We may not believe in a unique true input value for the model;
(ii) We may be unsure whether there are any good choices of input parameters (due to model deficiencies)
(iii) Full Bayes calibration analysis may be very difficult/non-robust.

A conceptually simple alternative is “history matching”, i.e. finding the collection of all input choices x for which you judge the match of the model to the data, $\|z - f_h(x)\|$ to be acceptably small, using some ”‘implausibility measure”’ $I(x)$ based on a natural probabilistic metric, accounting for emulator uncertainty, condition uncertain, structural discrepancy, observational error etc.

History matching via Implausibility

Using the emulator we can obtain, for each set of inputs x, the mean and variance, $\mathrm{E}(f_h(x))$ and $\mathrm{Var}(f_h(x))$.

History matching via Implausibility

Using the emulator we can obtain, for each set of inputs x, the mean and variance, $\mathrm{E}(f_h(x))$ and $\mathrm{Var}(f_h(x))$.

As $z_i = y_i + e_i$, $y_i = f_i(x^*) + \epsilon_i$,
if $x = x^*$, then
$\mathrm{Var}(z_i - \mathrm{E}(f_i(x))) = \mathrm{Var}(f_i(x)) + \mathrm{Var}(\epsilon_i) + \mathrm{Var}(e_i)$.

History matching via Implausibility

Using the emulator we can obtain, for each set of inputs x, the mean and variance, $\mathrm{E}(f_h(x))$ and $\mathrm{Var}(f_h(x))$.

As $z_i = y_i + e_i$, $y_i = f_i(x^*) + \epsilon_i$,
if $x = x^*$, then
$\mathrm{Var}(z_i - \mathrm{E}(f_i(x))) = \mathrm{Var}(f_i(x)) + \mathrm{Var}(\epsilon_i) + \mathrm{Var}(e_i)$.

We can therefore calculate, for each output $f_i(x)$, the "implausibility" if we consider the value x to be the best choice x^*, which is the standardised distance between z_i and $\mathrm{E}(f_i(x))$, which is

$$I_{(i)}(x) = |z_i - \mathrm{E}(f_i(x))|^2 / [\mathrm{Var}(f_i(x)) + \mathrm{Var}(\epsilon_i) + \mathrm{Var}(e_i)]$$

[Large values of $I_{(i)}(x)$ suggest that it is implausible that $x = x^*$.]

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

With this information, we can then refocus our analysis on the 'non-implausible' regions of the input space, by

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

With this information, we can then refocus our analysis on the 'non-implausible' regions of the input space, by

(i) making more simulator runs

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

With this information, we can then refocus our analysis on the 'non-implausible' regions of the input space, by

(i) making more simulator runs

(ii) refitting our current emulators

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

With this information, we can then refocus our analysis on the 'non-implausible' regions of the input space, by

(i) making more simulator runs

(ii) refitting our current emulators

(iii) emulating additional outputs (which were hard to emulate in the original parameter space)

over such sub-regions and repeating the analysis.

Using Implausibility measures

The implausibility calculation can be performed univariately, or by multivariate calculation over sub-vectors. The implausibilities are then combined, such as by using $I_M(x) = \max_i I_{(i)}(x)$, and can then be used to identify regions of x with large $I_M(x)$ as implausible, i.e. unlikely to be good choices for x^*.

With this information, we can then refocus our analysis on the 'non-implausible' regions of the input space, by

(i) making more simulator runs

(ii) refitting our current emulators

(iii) emulating additional outputs (which were hard to emulate in the original parameter space)

over such sub-regions and repeating the analysis.

This process is a form of iterative global search aimed at finding all choices of x^* which would give good fits to historical data.

Forecasting

The mean and variance of $f(x)$ are obtained from the mean function and variance function of the emulator f for F. Using these values, we compute the mean and variance of $f^* = f(x^*)$ by first conditioning on x^* and then integrating out x^* (over the space remaining after the history match).

Forecasting

The mean and variance of $f(x)$ are obtained from the mean function and variance function of the emulator f for F. Using these values, we compute the mean and variance of $f^* = f(x^*)$ by first conditioning on x^* and then integrating out x^* (over the space remaining after the history match).
Given $\mathrm{E}(f^*)$, $\mathrm{Var}(f^*)$, and the model discrepancy, ϵ and sampling error e variances, it is now straightforward to compute the joint mean and variance of the collection (y, z) (as $y = f^* + \epsilon, z = y_h + e$).

Forecasting

The mean and variance of $f(x)$ are obtained from the mean function and variance function of the emulator f for F. Using these values, we compute the mean and variance of $f^* = f(x^*)$ by first conditioning on x^* and then integrating out x^* (over the space remaining after the history match).

Given $\mathrm{E}(f^*)$, $\mathrm{Var}(f^*)$, and the model discrepancy, ϵ and sampling error e variances, it is now straightforward to compute the joint mean and variance of the collection (y, z) (as $y = f^* + \epsilon, z = y_h + e$).

We now evaluate the adjusted mean and variance for y_p adjusted by z using the Bayes linear adjustment formulae. This analysis is tractable even for real time control of large systems under complex forms of reification.

Forecasting

The mean and variance of $f(x)$ are obtained from the mean function and variance function of the emulator f for F. Using these values, we compute the mean and variance of $f^* = f(x^*)$ by first conditioning on x^* and then integrating out x^* (over the space remaining after the history match).

Given $\mathrm{E}(f^*)$, $\mathrm{Var}(f^*)$, and the model discrepancy, ϵ and sampling error e variances, it is now straightforward to compute the joint mean and variance of the collection (y, z) (as $y = f^* + \epsilon, z = y_h + e$).

We now evaluate the adjusted mean and variance for y_p adjusted by z using the Bayes linear adjustment formulae. This analysis is tractable even for real time control of large systems under complex forms of reification.

(When the forecast variance is large, then we have methods to improve forecast accuracy.)

Concluding comments

To assess our uncertainty about complex systems, it is enormously helpful to have an overall (Bayesian) framework to unify all of the sources of uncertainty.

Concluding comments

To assess our uncertainty about complex systems, it is enormously helpful to have an overall (Bayesian) framework to unify all of the sources of uncertainty.

Within this framework, all of the scientific, technical, computational, statistical and foundational issues can be addressed in principle.

Concluding comments

To assess our uncertainty about complex systems, it is enormously helpful to have an overall (Bayesian) framework to unify all of the sources of uncertainty.

Within this framework, all of the scientific, technical, computational, statistical and foundational issues can be addressed in principle.

Such analysis poses serious challenges, but they are no harder than all of the other modelling, computational and observational challenges involved with studying complex systems.

Concluding comments

To assess our uncertainty about complex systems, it is enormously helpful to have an overall (Bayesian) framework to unify all of the sources of uncertainty.

Within this framework, all of the scientific, technical, computational, statistical and foundational issues can be addressed in principle.

Such analysis poses serious challenges, but they are no harder than all of the other modelling, computational and observational challenges involved with studying complex systems.

In particular,
Bayes (linear) multivariate, multi-level, multi-model emulation,
careful structural discrepancy modelling
iterative history matching and forecasting
gives a great first pass treatment for most large modelling problems.

Concluding comments

To assess our uncertainty about complex systems, it is enormously helpful to have an overall (Bayesian) framework to unify all of the sources of uncertainty.

Within this framework, all of the scientific, technical, computational, statistical and foundational issues can be addressed in principle.

Such analysis poses serious challenges, but they are no harder than all of the other modelling, computational and observational challenges involved with studying complex systems.

In particular,
Bayes (linear) multivariate, multi-level, multi-model emulation,
careful structural discrepancy modelling
iterative history matching and forecasting
gives a great first pass treatment for most large modelling problems.
Great resource: the Managing Uncertainty in Complex Models web-site
http://www.mucm.ac.uk/ (for references, papers, toolkit, etc.)
[MUCM is a consortium of U. of Aston, Durham, LSE, Sheffield, Southampton - with Basic Technology funding. Now mutated into the MUCM community.]

References

P.S. Craig, M. Goldstein, A.H. Seheult, J.A. Smith (1997). Pressure matching for hydocarbon reservoirs: a case study in the use of Bayes linear strategies for large computer experiments (with discussion), in Case Studies in Bayesian Statistics, vol. III, eds. C. Gastonis et al. 37-93. Springer-Verlag.

J. Cumming, and M. Goldstein (2009). Bayes Linear Uncertainty Analysis for Oil Reservoirs Based on Multiscale Computer Experiments. In The Oxford Handbook of Applied Bayesian Analysis. O'Hagan, and West (eds), Oxford University Press. 241-270.

I. Vernon, M.Goldstein, and R. Bower (2010) Galaxy Formation: a Bayesian Uncertainty Analysis (with discussion), Bayesian Analysis, 5, 619-670

D. Williamson, M. Goldstein, et al (2013) History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. Climate Dynamics, 41 (7-8). 1703-1729

M. Goldstein and J.C.Rougier (2008). Reified Bayesian modelling and inference for physical systems (with discussion), JSPI, 139, , 1221-1239

M. Goldstein (2006) Subjective Bayesian analysis: principles and practice (with discussion) Bayesian Analysis 1 403-420