

Computer Predictions with Quantified Uncertainty, Part I

By Tinsley Oden, Robert Moser, and Omar Ghattas

Recently, a fresh look at how computer predictions are made has combined with several old philosophical ideas to bring about a revolution in computational science. The resulting panorama of formidable new challenges and research opportunities have to do with what computer models were always intended to do: make predictions of physical reality. Today, however, the phenomena and processes we ask computer models to predict are of enormous importance to critical decisions that affect our welfare and security—concerning, for example, climate change, the performance of energy and defense systems, the biology of diseases, and the outcome of medical procedures. With such high stakes, we must insist that the predictions include concrete, quantifiable measures of uncertainty. In other words, we must know how good the predictions are. The term “predictive simulation” has thus taken on a special meaning: the systematic treatment of model and data uncertainties and their propagation through a computational model to produce predictions of quantities of interest *with quantified uncertainty*.

The Predictive Science Process

Our goal is to use scientifically based predictions of physical reality to make informed decisions. Scientifically based predictions are simply forecasts of physical events and processes based on the methods of science: scientific theories (assertions about the underlying reality that brings about a phenomenon) and observations (knowledge received through the senses or the use of instruments).*

For our purposes, theory and observation—the fundamental pillars of science—can be cast as mathematical models: mathematical constructs that describe a system and represent knowledge of the system in a usable form. Mathematical models are thus abstractions of physical reality. Fundamental to all of science, they have been used successfully for millennia. However, mathematical models generally involve parameters that must be “tuned” so that the model best represents the particular system or phenomenon about which predictions are to be made. These are the moduli, coefficients, solution domains, boundary and initial data, etc., that distinguish one model from another within a class determined by the theory selected to characterize the physical phenomenon. Unfortunately, these model parameters are commonly not known with great precision; they may vary from material to material, specimen to specimen, and case to case, or they may not be known at all. In short, they generally involve large uncertainties that can be resolved only with sufficient experimental evidence. At the same time, experimental observations themselves are often fraught with errors and uncertainties owing to imperfections in the instruments or the difficulty or impossibility of acquiring observational data relevant to the problem at hand.

The digital computer, one of the greatest triumphs of modern science and technology, makes possible the use of mathematical models of enormous complexity—leading to the recognition of computer modeling and simulation as the third indispensable pillar of science. But this comes at a cost. Mathematical models are often corrupted as we create the computational models that render them amenable to solution via computer, and this corruption introduces more errors. Thus, a cascade of errors and uncertainties infect every aspect of scientifically based predictions (see Figure 1).

How can we cope with these imperfections? It is here that old ideas—from philosopher Karl Popper (1902–1994) and theologian and mathematician Thomas Bayes (1702–1761)—re-emerge.

In the Popperian view of scientific philosophy, a hypothesis did not qualify as a legitimate scientific theory unless it could be refuted by experimental evidence. This is the *principle of falsification*, the answer to the *problem of induction*, a great paradox in scientific philosophy from the time of David Hume in the 18th century. In this view, a scientific theory can never be validated; it can only be invalidated by contrary experimental evidence. Experimental observations, then, are intrinsically interwoven into the scientific method: Without a possible program of experiments, scientific theo-

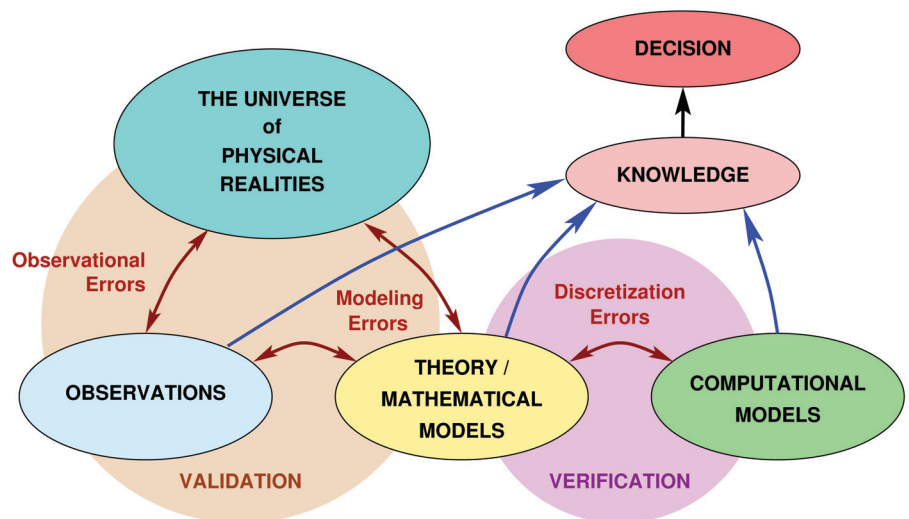


Figure 1. Imperfect computational modeling: Imperfections in the mathematical models, incomplete observational data, observations delivered by imperfect instruments, and corruption of the model itself in the discretization needed for computation all lead to imperfect paths to knowledge. Reproduced from J.T. Oden, “A Brief View of V & V & UQ,” a presentation to the Board on Mathematical Sciences and Their Applications, National Research Council, October 2009.

*This subject is the focus of the National Nuclear Security Administration’s Predictive Science Academic Alliance Program. The mathematical foundations of verification, validation, and uncertainty quantification are also the focus of a new study organized by the Board on Mathematical Sciences and Their Applications of the National Research Council.

ries are not legitimate. Analogously, we need experimental observations in order to test the validity of the mathematical models we use in computer simulations.

But the Popperian view has been criticized for being a canon of *objective philosophy*, which is regarded by many as too rigid, not reflective of the way science is really done. Indeed, theories and their consistency with experiments, it is argued, must be judged in terms of probabilities in light of the evidence. This softer interpretation of validity fits more closely with the philosophy of Thomas Bayes, from whose writings grew Bayesian statistics and the simple but powerful Bayes theorem. While laid down two and a half centuries ago, the Bayesian view of the scientific method and its role in prediction has been recognized only recently as a great unifying framework that elegantly ties all the components of predictivity together.¹ Before discussing this framework, we consider several important *processes* of predictive science:

Identification of Quantities of Interest. An idea not widely appreciated by computer modelers only a few years ago is that the entire exercise of simulating a physical phenomenon begins with a clear specification of the goals of the calculation: What are the target outputs? What are the principal “quantities of interest,” or QoI’s? At the heart of the verification and validation processes discussed below is the realization that a model might be perfectly acceptable for simulating some features of a phenomenon, but entirely inadequate for modeling other features. Before asking whether a model is or is not invalid, we must specify the quantities we wish to predict, and the situations or scenarios in which they are to be predicted.

Clearly, the quantities of interest for a predictive simulation are dictated by the larger purpose of the prediction—to assess, for example, the safety of an aircraft design or the chance that a medical procedure will be successful. Often, such QoI’s appear mathematically as numbers, that is, as functionals on the solutions of the governing equations. In any case, the verification and validation processes, and indeed the prediction itself, must be built around the calculation of specific QoI’s.

Verification. Because computational models are corrupted (approximate) versions of the mathematical models on which they are based, we must confirm that a computation accurately approximates results for the QoI’s produced by the underlying mathematical model. This is verification, the process designed to detect and control errors brought about by discretization of the mathematical model and bugs introduced in implementation of the model in software.

Calibration. In adjusting parameters in mathematical models to bring them into (closer) agreement with experimental observations, responses of components of a system are measured. The corresponding model parameters are then inferred from the measurements in what amounts to an inverse problem (given the observed outputs, find the corresponding input parameters). Such inverse problems are commonly ill-posed or under-determined, requiring that the modeler provide further information regarding the nature of the parameters and/or the quality of the expected correspondence between model and observations. As discussed below, when considering uncertainty, this can be done through Bayesian inference.

Validation. The validation process is designed to build confidence that the mathematical model can accurately predict the QoI’s. Thus, validation addresses the question of the legitimacy of the theoretical model for the purposes of the predictions to be made. Validation processes necessarily involve a carefully designed program of physical experiments that are intended to assess the degree to which the model can reproduce experimental observations. But the process is complicated by the fact that the QoI’s are commonly not accessible for observation. To determine whether a set of experimental observations invalidate a model for the purposes of a particular prediction, the modeler must thus consider the degree of disagreement between model and observations, the impact of these disagreements on predictions of the QoI’s, and the tolerance for error in the predictions.

Use of the word “process” in the definitions of verification and validation is important: It is generally impossible to completely verify a computational model, and, in line with Popper’s principle, a model can never actually be validated.

Uncertainty Formulations

The preceding descriptions of predictive simulation processes are independent of considerations of uncertainty; how these processes are actually conducted, however, depends critically on the treatment of uncertainty. In treating uncertainty, we must first decide how to represent it mathematically. A number of representations have been proposed, based on such formalisms as the Dempster–Shafer theory of evidence, the theory of fuzzy sets, interval analysis, worst-case scenarios, and many more. The Bayesian philosophy discussed earlier suggests that uncertainty be represented probabilistically, a particularly powerful approach that we discuss in some detail below. The underlying concepts are more general, however, and they are relevant regardless of the uncertainty formalism used.

Probabilistic representation of uncertainty requires some explanation. Probability most commonly represents random processes, and in computational predictions uncertainties of one type indeed arise from randomness. These are *aleatoric* uncertainties, which are irreducible, in that better data or improved models cannot reduce them. In many other cases, the dominant uncertainties arise from lack of knowledge, particularly lack of knowledge of physical model parameters and imperfections in the mathematical models themselves. These are *epistemic* uncertainties, which can in principle be reduced. In the Bayesian formulation, epistemic uncertainties are also represented probabilistically. In this case, probability represents our confidence in some proposition, given all currently available information. Fundamentally, the Bayes theorem gives us the formalism to update that confidence (probability) when new information becomes available.

In the Bayesian uncertainty formulation, we make no distinction between epistemic and aleatoric uncertainties. Instead, this distinction is one of interpretation, in the context of the prediction to be made, of what information could be obtained on which to base the prediction. An example can be found in an epitome of random processes, the throwing of dice (“aleatoric” is from the Latin *alea*, meaning dice game or die). In the context of games of chance, the outcome of a throw of fair dice is considered random. The physical processes controlling this outcome, however, are themselves deterministic. All that is lacking for a reliable prediction is detailed knowledge of the initial conditions of the throw; this information is interpreted

¹See C. Howson and P. Urbach, *Scientific Reasoning: The Bayesian Approach*, 2006, for a modern account. The view that scientific theories must be judged in terms of their probabilities was shared by many of the foremost scientific minds of history, including Laplace (1749–1827), Poincaré (1854–1912), and Jaynes (1922–1998).

as unattainable at a casino, where a dice game is, appropriately, ruled by aleatoric uncertainty. But in another context, perhaps in a lab with a dice-throwing machine, it might be possible to measure or control these initial conditions. In this case, imperfect knowledge of the initial conditions makes the outcome of a throw epistemically uncertain.

Despite the indifference of Bayesian analysis to aleatoric versus epistemic uncertainties, the distinction is nonetheless useful. Consider the example of simulations intended to predict the failure of widgets built at a widget factory. If the uncertainty in the simulations is primarily aleatoric (e.g., due to random variations in the dimensions of the parts), then the predicted probability of failure is the predicted percentage of widgets from the factory that will fail. But if the uncertainty is primarily epistemic (e.g., from imperfections in the computational model), then the widgets produced at the factory will presumably all fail or all work; the failure probability characterizes our confidence in predicting which will happen. Clearly, the difference in these two scenarios would be important to the warranty claims office at the factory.

Calibration and Statistical Inverse Problems

In the Bayesian analysis of computational predictions, uncertainty arises from three sources, each, as discussed earlier, represented probabilistically[†]:

Model parameters. Instead of a list of parameters m , we begin with the prior joint probability density function (pdf) $\rho_M(m)$ describing what we know about the parameters before we begin the modeling exercise. These parameter distributions are gleaned from the literature, previous experiments, or archival databases, or specified to represent almost complete ignorance of the parameters. Calibration via Bayesian inference will be used to update the distributions to make the theory (model) consistent with particular observations (data).

Experimental observations. As noted earlier, the observational data d also will have uncertainties, also represented by pdf's, $\rho_D(d)$.

Theoretical model. The mathematical model used to make the predictions is a key component of the overall process—it defines the way given parameter sets m are mapped into the theoretical observations d . But because the theory (model) may not be completely reliable, our uncertainty or lack of confidence in the theory must be included in the analysis. We accomplish this by expressing the theory as a conditional probability distribution that maps the model inputs to a probability distribution of the outputs. The conditional probability $\theta(d|m)$ (“ θ ” indicating “theory”) is the probability that the model will predict the observations d given the parameters m . This is also known as the *likelihood*.

Various forms of Bayes’s theorem pull these elements together to characterize exactly what is known about the model parameters, expressed as the posterior pdf $\sigma(m|d)$ of the parameters m conditioned on the data d . From the theorem,

$$\sigma(m|d) = \frac{\rho_M(m)\theta(d|m)}{\rho_D(d)}. \quad (1)$$

This expression is essentially a statistical calibration, as it statistically infers the (posterior) distribution of parameters $\sigma(m|d)$ that fit the theoretical model to the observations d .

Bayes’s theorem and the statistical calibration process embodied in (1) are also the keys to validation and, ultimately, prediction with quantifiable uncertainties. These topics will be discussed in Part II of this article, which will appear in an upcoming issue of *SIAM News*.

Acknowledgments

We have benefited from numerous discussions with I. Babuška and R. Tempone and from their published work on verification, validation, and uncertainty quantification. We thank the staff of the PECOS Center for their contributions to pursuing these ideas, and in particular T. Oliver, E. Prudencio, and S.H. Cheung for the results shown in Figure 4 (to appear in Part II). The financial support of the PECOS Center by the Department of Energy under the National Nuclear Security Administration’s Predictive Science Academic Alliance Program is gratefully acknowledged (DE-FC52-08NA28615).

Tinsley Oden, Robert Moser, and Omar Ghattas lead the Center for Predictive Engineering and Computational Sciences (PECOS) of the Institute for Computational Engineering and Sciences at the University of Texas at Austin.

[†]See, for example, J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005.