In *Proc. 8th Nuclear Explosives Code Developers Conference*, E. Caramana and C. McMillan, eds., pp. 457-462, LANL Report LA-12963-C (1995)

LA-UR-01-6671
LA-CP-94-0288

Validation of Hydrocode Predictions

Kenneth M. Hanson and Gregory S. Cunningham
Los Alamos National Laboratory (Dynamic Experimentation Division)

A new strategy for validation of hydrocode predictions through the Bayesian analysis of radiographic data is presented. The Bayesian approach allows one to incorporate prior knowledge about the structure of the objects being analyzed and provides the foundation for assessing the reliability of the results. We propose using the hydrocode prediction at radiographic time as the initial or default object model. The object model is altered from its default in a minimal way to match the available radiographs in the Bayesian sense. A full understanding of the degree of validity of the final model relies on the ability to explore and characterize the uncertainty in the model, a relatively new feature in Bayesian analysis. We suggest that a physics-based validation of hydrocodes themselves may require some, if not all, of the basic concepts presented here to infer aspects of the underlying physics models from hydrodynamic experiments.

Introduction

Our present approach to benchmarking hydrodynamical simulations is based on a direct comparison between the 3D object predicted by the hydrocode at radiographic time and the best reconstruction that can be obtained from the available radiographic data. This approach is fraught with difficulties when the number of experiments that can be performed is limited and the object does not possess significant symmetry. With our present facilities, we can obtain only one radiograph in each dynamic experiment. When the object under study possesses axial symmetry, one radiograph is sufficient to reconstruct the object. We have demonstrated that we can analyze radiographs of objects that nearly possess axial symmetry very accurately, in terms of locating boundaries, estimating densities, and tomographic reconstruction.

On the other hand, a completely unambiguous reconstruction of a complex 3D object requires many radiographs. For example, to fully resolve the ambiguities at a one-mm resolution for a 10-cm-wide 3D object with an arbitrary density distribution, approximately 100 radiographs are necessary. Fortunately, many of the dynamic objects with which we deal do not possess completely arbitrary density distributions. Their density distributions are often relatively slowly varying in some regions and various symmetries may often be assumed. As a consequence, we can often obtain a reasonably good reconstruction of a dynamic object from a handful of radiographs (Mathews, 1994).

In practice, to reconstruct a dynamic object we repeat the experiment a number of times with the device placed in different orientations relative to the radiographic axis. However, the high cost of each experiment limits the number of experiments that can be done to a handful. Furthermore, the reproducibility of the details of the dynamic object between experiments must be assumed. Therefore, we are placed in the position of trying to compare a radiographic reconstruction, which we know is imperfect for lack of sufficient data, to a hydrocode prediction. When we see differences between such a reconstruction and a hydrocode prediction, how can we infer what is valid and what is invalid about the hydrocode result?

To gain as much information as possible about the validity of a hydrocode prediction from whatever experimental data are available, we propose to make more direct use of the hydrocode prediction. We suggest that the 3D object predicted by the hydrocode be used as a starting point for the analysis of the experimental data. The shape of the 3D object can be deformed and its densities adjusted to make its predicted radiographs match the experimental ones. An analysis of the reliability of the final reconstruction guides the conclusions that can be drawn from the experiments. The Bayesian methodology provides suitable means to do this.

Summary of the Bayesian Approach

In this short paper, we do not have the space to thoroughly develop the Bayesian approach. The reader is urged to refer to (Hanson, 1987, Hanson, 1993a, Hanson and Cunningham, 1994) for more detail.

The Bayesian approach is based on the use of parametric models to describe the object of interest. In Bayesian analysis uncertainties in parameter values are represented by probability distributions on those parameters. A relatively large uncertainty in a parameter is represented by a broad distribution; a precisely known parameter by a narrow distribution. Probability theory provides a quantitative and consistent basis for the Bayesian analysis, which inherits its name from Bayes fundamental law governing the updating of one probability distribution, called the prior, in the face of new data, called the likelihood, to obtain the resulting probability, called the posterior.

The essential action of Bayes law is captured in the theory of the propagation of experimental errors to which most scientists are exposed early in their studies. When accurate measurements are combined with less accurate ones, coming from prior experiments for example, the uncertainty in the combined result will be significantly reduced compared to that before the acquisition of the new data. By providing a much more thorough description of uncertainty in the form of a precise probability distribution, Bayesian analysis allows one to address more detailed issues, for example, optimal estimators, confidence intervals in the estimates (for arbitrary probability distributions), etc.

In addition to estimation of the uncertainties in parameter values, the Bayesian approach also provides a methodology for inference from measurements about the choice of models appropriate to describe reality. Probabilities can ultimately be employed to compare and decide amongst different models. Our preference for simpler models over more complex ones can be incorporated through a prior on model complexity (Gull, 1989). A simple example of where model selection plays a role is in the choice of the number of terms to use in a series expansion, which is meant to describe the behavior of some data.

Bayesian analysis also provides the means to properly make subsequent decisions through the use of cost or utility functions, which specify the costs of making correct versus incorrect decisions. Examples of such kinds of decisions include, in the field of nondestructive testing, whether to accept or reject a precision part on the basis of a radiograph, or, in medicine, whether to follow up a possibly positive indicator in a screening test with another test or surgery.

The full state of our knowledge about reality is summarized by the posterior probability, or simply the posterior. The standard approach to obtaining a representative solution is to find the parameter values that maximize the posterior, called the MAP solution. Although this single solution is often the goal for many investigators, the posterior probability can be more fully utilized to determine the degree and character of the uncertainty in the solution.

We are developing a tool for exploring the posterior to provide an understanding of the degree of uncertainty in Bayesian solutions, which we describe below.

Bayes Inference Engine

We are developing the Bayes Inference Engine (BIE) to implement the Bayesian methodology on a computer workstation. Our goals for the BIE are that it should be easy to learn and use, and that it should provide a high degree of interactivity with good visualization of the process and the models. Additionally, we intend to build an application that provides the user with a great deal of flexibility in configuring object models and measurement models. We deem these features essential to the usefulness of the BIE.

In the following discussion we will use the symbol ϕ for minus the logarithm of the posterior probability. Computations with this function are typically easier to do than with the posterior itself since the products of probabilities in Bayes law become sums. The most probable parameter values in a MAP solution then occur at a minimum in ϕ.

Brief Overview of the BIE

The BIE is programmed in the object-oriented language Smalltalk in the version supplied by ParcPlace Systems[1], which includes a complete class library for user-interface development. The interface to the BIE is the graphical programming tool (Cunningham et al. 1994a), which operates as follows (refer to Fig. 1). One is presented with a canvas, on which appear buttons that allow the user to add items to, or delete items from, the canvas. One can add or delete Transforms and Connections. In this description the capitalized words are objects in the object-oriented language. Transforms act on input Data to calculate output Data and are represented on the screen by a square icon. One specifies the data-flow by connecting one Transform to another using a Connection, which is represented by lines drawn between the two Transforms.

The Transforms are living objects and one can interact with them in several ways. One can see a description of a Transform and change the parameters that define it. Through a menu that pops up when one clicks on a Transform icon, one can have the Transform display its output. The fact that the Transform objects are alive distinguishes this graphical programming tool from one that allows a user to construct and visualize a script that contains a sequence of actions to be executed (off line) in the prescribed order.

Referring to the data-flow diagram in Fig. 1, the Parameters of the object model (the leftmost icon) provide input to the measurement model. The radiographic measurement model shown consists of the next three icons, which sequentially take the projection of the object, exponentiate the result, and perform a

[1] 900 East Arques Ave., Sunnyvale, CA 94086

blurring to mimic radiographic blur. The output of the measurement model is predicted data, which are fed into a LogLikelihood function, designated by ϕ, along with the actual data, the uppermost icon. A LogPrior, which operates on the model parameters, can also be specified. The LogPrior + LogLikelihood feed the LogPosterior, which can be optimized with respect to object-model parameters using conjugate gradient (for unconstrained problems) and a modified gradient descent (for constrained problems) strategies. One specifies that the Parameters of the object model are to be optimized by connecting the Parameter icon to the Optimizer, the lower right-hand icon. After optimization, in which the minimum of the LogPosterior is found, the object model and its Parameter values represent the MAP solution.

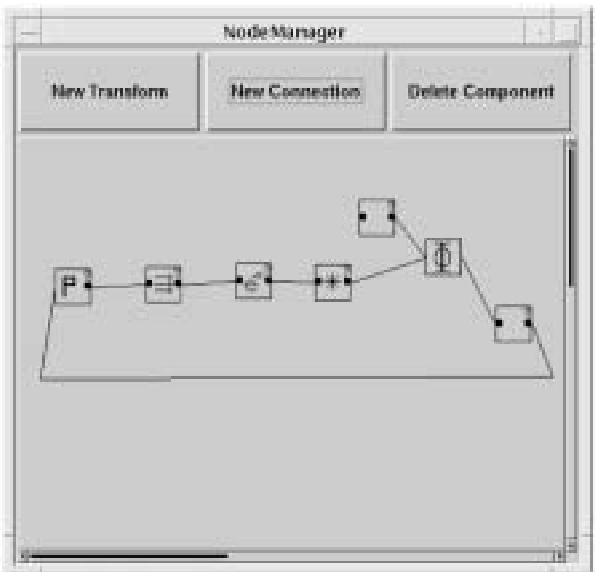

Figure 1. The canvas of the Bayes Inference Engine permits one to specify a data-flow diagram by connecting together Transforms.

The BIE incorporates many innovative features including: 1) geometrical representations of physical objects, 2) adjoint differentiation to calculate the gradient of ϕ with respect to all object parameters, 3) new approaches to solving the optimization problem, which is required to find the MAP solution, 4) a new method to explore the reliability of the solution, and 5) a graphical-programming interface based on object-oriented programming technology, which greatly enhances the flexibility of modeling objects and measurement processes. The details concerning these new developments are to be found in the references. However, several of these innovations are worth mentioning in the context of this conference proceedings.

Geometric Representation of Objects

We are pioneering the concept of describing physical objects in terms of their geometry to improve tomographic reconstruction. This tack is quite different from the normal one of representing a 3D object in terms of its density, typically described by cubical voxels on an ordered grid. The use of a geometrical description recognizes the very important role that boundaries play in characterizing objects. The reconstruction process amounts to deforming an initial object geometry in a minimal way to match the data. In the Bayesian approach, one controls the geometric deformation by placing a prior on it. The net effect is to add an energy of deformation to ϕ so that greater deformations are penalized in the reconstruction procedure (Hanson, 1993b). This approach has proven to be a valuable means to achieve good reconstructions in situations where all other methods fail, for example when only a few radiographs are available. However, it must be emphasized that this approach can only be successful when the objects being reconstructed have a fairly simple morphology that is approximately known beforehand. We intend to use hydrocode predictions to provide this kind of shape information to the reconstruction procedure.

Adjoint Differentiation

We have uncovered a little-known technique called adjoint differentiation (Thacker, 1991). In our application we need to minimize the scalar function ϕ by varying the many (10^3 to 10^6 or more) variables that comprise the parameters of the object model. This optimization problem would be insoluble without knowing the gradient of ϕ, or sensitivities, with respect to the many parameters on which it depends. The adjoint differentiation technique facilitates this calculation in a computational time that is comparable to the forward calculation through the data-flow diagram. Our use of objects to represent transforms greatly aids the implementation of this adjoint calculation (Cunningham et al., 1994b).

The adjoint differentiation technique we are advocating is closely related to the differential sensitivity approach being pursued in T and X Divisions by Maudlin et al. (1993). Their approach is based on the dynamic equations for the adjoints to the sensitivities of the real physical quantities. They solve these equations by standard integration techniques (presently based on MESA), working backwards in time. When solving nonlinear differential equations, this approach must make reference to the complete forward solution to the governing physics equations. Our approach differs in that it provides the adjoint sensitivity calculation that exactly matches each specific forward updating step.

Reliability Exploration

Another innovation that we wish to mention here, because of its close connection to physics and the central role it can play in understanding the implications of a Bayesian analysis, is that of reliability exploration. There exist at least two ways to visualize the reliability of inferred plausible models. The first, proposed by Skilling et al. (1991), provides a stochastic look at the range of possible solutions. It involves the display of a sequence of solutions that are randomly chosen from the posterior probability distribution. This sequence, typically calculated off line, is presented as a video loop. By showing a representative range of alternative solutions, the degree of variability of this presentation provides the viewer with a graphic impression of the degree of uncertainty in the inferred model.

Our new approach (Hanson and Cunningham, 1994) draws on an analogy between ϕ and a physical potential. Then the gradient of ϕ is analogous to a force. An unconstrained MAP solution can be interpreted as the situation in which the forces on all the variables in the problem balance so that the net force on each variable is zero. Furthermore, when a variable is perturbed from the MAP solution, the derivative of ϕ with respect to that variable is the force that drives it back towards the MAP solution. The phrase force of the data takes on real meaning in this context.

We propose to exploit this physical analogy to facilitate the exploration of the reliability of a particular feature of a MAP solution, which the user specifies by directly interacting with the solution presented by the BIE. The uncertainty in the solution is explored by applying a constant force to the selected combination of parameters that characterize the feature of interest. All parameters are readjusted to minimize ϕ. The uncertainty in the combination of parameters is indicated by the rate at which they move away from their MAP value as the external force is applied to them. The correlations between parameters experiencing the external force and the others are demonstrated by how much and in what direction the parameters change. Ideally, these correlations could be seen through direct interaction with a rapidly-responding dynamical Bayesian system. Alternatively, they may be demonstrated by means of a video loop.

Another interesting aspect of this technique is the possibility of decomposing the forces into components. For example, the force derived from all data (through the likelihood), or even a selected set of data, may be compared to the force derived from the prior.

We anticipate that it may be possible in the future to use the tools of virtual reality, coupled to turbocomputation, to explore the reliability of a Bayesian solution through direct manipulation of the computer model. Force feedback would permit one to actually feel the stiffness of a model. Higher dimensional correlations might be felt through one's various senses.

Validation of Hydrocode Predictions

Consider the following scenario: a single radiograph of a complex dynamic object is available. Clearly tomographic reconstruction in the normal sense is impossible. How is one to proceed with the task of validating a hydro prediction? From the hydrocode prediction one can calculate what the radiograph should look like. So it is possible to compare the predicted radiograph with the actual one. Suppose that, to within the known experimental uncertainty in the measured radiograph, these two radiographs match each other down to the last detail. Then clearly the hydrocode prediction has been validated, at least in regard to those aspects of the predicted object that can be determined by that unique radiograph. Of course, the densities at many points in the predicted object contribute to each pixel in the observed radiograph. Thus, when a single or even several radiographs are properly predicted by a hydrocode prediction, what has been learned is that a certain combination of densities of the predicted object is correct, to within the accuracy of the measurements. In other words, not all aspects of the hydrocode are validated from just a few radiographs.

To make the most of meager data, we propose to use the hydrocode prediction at radiographic time as a starting point for the analysis of a set of hydrodynamic radiographs. The radiographic data provide the basis for the inference process. If they are not matched by the data predicted by the object model, then the object model should be changed, by deforming the geometry of the object or adjusting the density distribution within the object. If the data can be matched by slightly molding the hydrocode prediction, then we judge how much the hydrocode prediction has missed the mark by how much it had to be altered. Of course, it must be understood that all inferences about the object have associated with them some degree of uncertainty, which we discuss next. If the hydrocode prediction turns out not to be a credible model of the actual object, an alternative choice for the starting object model may required. If essential features of the hydrocode prediction are not observed in the radiographic data, the inference is that the hydrocode prediction does not properly predict reality and needs refinement.

The aim of any inference process is to determine those aspects of a model that are determined well by an experiment or experiments and those that are not. The science (or art?) of experimental design is to select experiments that reveal the most about those aspects of the models that need to be known. A good experiment gives detailed and accurate information about those critical features of the physical model that are crucial to know without interference from other unknown features. Since we may not be able to guarantee that a dynamic object will be the same in repeated experiments, and because of limited money and time to do the experiments, we must learn to cope with a limited number of radiographs and make the most of them.

The purpose of this paper is to indicate that we are making headway in understanding how to ascertain the range of validity of hydrocode predictions in this scenario. However, this process will not be as simple as the old methodology of comparing a single 3D reconstruction from a series of radiographs to a hydrocode prediction and concluding that the prediction is either invalid or precisely valid.

Reliability of Derived Quantities

The reconstruction of an object is often used to calculate a derived quantity, for example, the reactivity implied by a reconstructed configuration of simulated nuclear material (Mathews, 1994). Since the Bayesian approach assigns a probability to every feasible configuration of the object, one can in principle generate the probability distribution of the derived quantity by calculating the derived quantity for every configuration. An approximation to such a probability distribution may be had by randomly drawing samples from the posterior probability distribution of the inferred object model, using the technique proposed by Skilling et al. (1991), or one similar to it, and calculating the derived quantity for each sample. The resulting frequency histogram represents the degree of the certainty one has in the derived quantity, which can be summarized by its estimated standard deviation from the mean, for example.

This approach may be useful in estimating the degree of certainty that we have in the values of reactivity calculated from radiographic data. In the spirit of the previous section, we may one day be able to estimate the uncertainty in the time dependence of the reactivity derived from a hydrocode calculation, after the code has been thoroughly validated through dynamic experimentation.

Implications for Validation of Hydrocodes

Although the first use of the concepts presented here will be employed in the BIE to validate the predictions of hydrocodes, we feel that many of the innovations described above could be beneficially employed in the validation of the hydrocodes themselves. Indeed, the aims of a physics-based approach to device assurance must be based on fully exploring, understanding, and quantifying our certainty of the various physical foundations on which our hydrocodes are based. This type of program would possess many of the same characteristics that are encompassed in the BIE. Therefore, the BIE might serve as a prototype for the approach that could be adopted in validating the hydrocodes themselves. Even if not every aspect of the BIE philosophy is followed, some of the ideas expressed here may be useful in hydrocode development.

We believe that if our adjoint differentiation technique were coupled to a similar one implemented within hydrocodes (Maudlin et al., 1993), it could be of great benefit, perhaps even essential, for accomplishing various tasks for which hydrocodes are used. One application of this technique could be in solving design problems with hydrocodes. If the design goals can be encapsulated in terms of a scalar objective function, which is to be minimized to meet the design criteria, the design problem then becomes an optimization problem. As in the BIE, the gradient of the objective function with respect to the variables in the problem, for example, the initial geometrical shape of key components, could be found. Then a general-purpose optimization algorithm, such as the one we are developing for the BIE, can be used to arrive at a good design.

Another area that could be greatly impacted is in addressing the validity of the models embedded in a hydrocode. The same methodology used in the BIE to draw inferences about models and model parameters can be applied to hydrocodes to help infer which aspects of its inherent physics models are well determined by experimental data, and which are not. The technique that we propose for exploring the reliability of a Bayesian interpretation of a particular model for reality could be used to directly explore the reliability of the physics modules contained in a hydrocode.

One of the strengths of the Bayesian approach is that it can easily synthesize the information from many data sets that bear on a particular physical model. Thus it should be possible to make inferences concerning a component of a hydrocode, for example, the equation-of-state of a particular material, on the basis of several different kinds of experiments that directly quantify the behavior of that material. This methodology could be extremely valuable to the designer by providing an understanding of the weaknesses in the physics models and where they affect hydrocode predictions. By extension, this approach can be used to design new experiments that would help fill in identified gaps in our knowledge.

Acknowledgments

We would like to acknowledge helpful conversations we have had with the following people: John Skilling, George Jennings, Jr., Paul Maudlin, Rudy Henninger, Richard Krajcik, Gary Wall, Harry Watanabe, Jay Boettner, Rollin Whitman, Bob Benjamin, Larry Hull, Eric Ferm, and Jack Jacobson.

References

Cunningham, G. S., Hanson, K. M., Jennings, Jr., G. R., and Wolf, D.R., "An Object-Oriented Implementation of a Graphical-Programming System," *Proc. SPIE*, **2167**, pp. 914-923 (1994a).

Cunningham, G. S., Hanson, K. M., Jennings, Jr., G. R., and Wolf, D. R., "An Object-Oriented Optimization System," *Proc. IEEE Int. Conf. Image Processing*, vol. III, pp. 826-830 (1994b).

Gull, S. F., "Bayesian Inductive Inference and Maximum Entropy," *Maximum Entropy and Bayesian Methods in Science and Engineering*, vol. 1, Erickson, G. J. and Smith, C. R., eds., (Kluwer Academic, Dordrecht, 1989), pp. 53-74.

Hanson, K. M., "Bayesian and Related Methods in Image Reconstruction from Incomplete data," in *Image Recovery: Theory and Application*, Henry Stark, ed., (Academic, Orlando, 1987), pp. 79-125.

Hanson, K. M., "Introduction to Bayesian image analysis," *Proc. SPIE*, **1898**, pp. 716-731 (1993a).

Hanson, K. M., "Bayesian Reconstruction Based on Flexible Prior Models," *J. Opt. Soc. Amer.*, **A 10**, pp. 997-1004 (1993b).

Hanson, K. M. and Cunningham, G. S., "The Hard Truth," in *Maximum Entropy and Bayesian Methods*, Skilling, J. and Sibisi, S., eds., (Kluwer Academic, Dordrecht, 1994), pp. 157-164.

Mathews, A. R., "Studies of One-Point Safety Using 3D Tomography (U)," *DX-13 Group Report*, DX-13-GR-94-2 (SRD).

Maudlin, P. J., Henninger, R. J., and Harstad, E. N., "Application of Differential Sensitivity Theory to Continuum Mechanics," *Proc. ASME Conf.*, New Orleans, 1993.

Skilling, J., Robinson, D. R. T., and Gull, S. F., "Probabilistic displays," in *Maximum Entropy and Bayesian Methods*, Grandy, Jr., W. T., and Schick, L. H., eds., (Kluwer Academic, Dordrecht, 1991), pp. 365-368.

Thacker, W. C., "Automatic Differentiation from an Oceanographer's Perspective," *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*, Griewank, A. and Corliss, F. C., eds. (SIAM, Philadelphia, 1991) pp. 191-201.