

An interactive tool for Bayesian inference

G. S. Cunningham, K. M. Hanson, G. R. Jennings, Jr., and D. R. Wolf

Los Alamos National Laboratory, MS P940
Los Alamos, New Mexico 87545 USA

ABSTRACT

The Bayes Inference Engine (BIE) is a flexible software tool that allows one to interactively define models of radiographic measurement systems and geometric models of experimental objects so that the geometric properties of the objects being radiographed can be inferred from a limited amount of data. The BIE also allows a user to investigate confidence intervals on the estimated object geometry and compare the likelihoods of competing hypotheses. The BIE contains three components: a graphical programmer for defining and interacting with the measurement system model, a geometric modeler for defining and interacting with the object model, and an interactive optimizer. This article contains a description of these three components and an example of 2D geometry optimization from synthesized radiographic data using the BIE.

Keywords: object-oriented modeling/programming, graphical programming, optimization, adjoint differentiation, Bayesian inference, tomographic reconstruction, geometric modeling

1. BACKGROUND

1.1. The Traditional Approach to Data Analysis

The traditional approach to data analysis starts with a “measurement model” H that describes how the data are obtained, on average, from a parameterization x of the object of interest. For example, Fig. 1 contains a simple radiographic measurement model that consists of a line integral transform A followed by exponentiation B and convolution C . Thus, given an object parameterization, one can calculate the data $y = H(x) = C(B(A(x)))$ that are predicted by the measurement model.

The measurement model can also be used to invert the data d to obtain an estimate of object parameters, $\hat{x} = H^{-1}(d)$. For example, one can deconvolve radiographic data, take the logarithm of the result, and finally perform an inverse Radon transform to obtain a pointwise estimate of the attenuation profile for a 2D object, $\hat{x} = A^{-1}(B^{-1}(C^{-1}(d)))$.

One problem with the traditional approach to data analysis is that the object parameterization must be such that the inverse of the measurement model is well-defined. For example, if only one radiographic projection of a 2D object is available, then a circularly-symmetric parameterization of the object must be used to invert the measurement model (using the inverse Abel transform).

Further author information:

G.S.C.: E-mail: cunning@lanl.gov

K.M.H.: E-mail: kmh@lanl.gov




Figure 1. A data-flow diagram for a simple measurement model for a radiographic system consisting of, from left to right, an image describing a density distribution, line integrals through that image, and negative exponentiation.

For complicated measurement models, though, it may not be easy to define a nice parameterization that allows the inverse to be well-defined.

Another problem with the traditional approach to data analysis is that no confidence intervals can be calculated to express the degree of certainty one should have in the estimated object parameters. That is, if x is a discrete attenuation profile in a 2D slice of an object and x_{ij} is the attenuation at a particular location, then the traditional approach to data analysis cannot answer a question like, “what is the probability that $a_1 \leq x_{ij} \leq a_2$?” Finally, the traditional approach does not allow one to evaluate competing hypotheses for explaining the data. That is, one cannot get answers to questions like “what is the likelihood that the attenuation profile of the object is constant in this region compared to the likelihood that it is constant with a flaw of this character?”

1.2. The Bayesian Approach to Data Analysis

The Bayesian approach to data analysis solves many of the problems associated with the traditional approach. The data vector predicted by the measurement model y is coupled to the measurement data d through a likelihood function $\phi(y, d) = \phi(H(x), d)$. $\phi(H(x), d)$ is the probability that the data obtained in an experiment would be equal to d given that the object has parameter values x . Thus, ϕ contains a complete description of the noise in the measurements. If the noise is additive, then $d = y + n$ and $\phi(H(x), d) = P_n(d - H(x))$, where $P_n(\cdot)$ is the probability distribution of the noise n .

Maximizing $\phi(H(x), d)$ over all permissible x given the data d yields a maximum-likelihood estimate (MLE) \hat{x} . If we have a prior probability distribution on x , say from a previous experiment or other measurements, then we can maximize the posterior probability distribution $\phi(H(x), d)\Pi(x)$ over all permissible x , given the data d , to produce a maximum *a posteriori* (MAP) estimate \hat{x} . In either case, \hat{x} contains the most likely values of the object parameterization given the data.

Note that the Bayesian approach only requires one to maximize $\phi(H(x), d)$ or $\phi(H(x), dd)\Pi(x)$, which does not require one to know H^{-1} . Furthermore, the Bayesian approach allows one to make probabilistic statements about the estimate \hat{x} . For example, one can ask “what is the probability that $a_1 \leq x_{ij} \leq a_2$?” or “what is the likelihood that the attenuation profile of the object is constant in this region compared to the likelihood that it is constant with a flaw of a specified character?”

The major problems associated with the Bayesian approach are the details of how to implement the optimization of the likelihood or posterior: 1) the likelihood and/or posterior may be highly nonlinear in the object parameterization x so that a global optimization is very difficult (there may be many local minima) 2) the dimension of x may be very large, making a gradient-based approach to optimization essential. Note that the latter problem also makes hypothesis testing more difficult, as one may find it difficult to formulate competing hypotheses that provide insight into the reliability of the estimates.

2. THE BAYESIAN INFERENCE ENGINE

2.1. BIE Design Goals

The BIE is our attempt to provide a software tool that solves the implementation problems associated with the Bayesian approach to data analysis.

There are two principles that guide us in developing the BIE. First, the tool should allow the user a high degree of graphical interaction with the measurement model and the object parameterization. This user interaction is necessary for easily controlling the complexity of the models (for global optimization) as well as for interrogating them in intuitive ways (for investigating the reliability of the estimate and intermediate predicted data). Second, the software should be written in an object-oriented (OO) language to maximize our productivity and provide the foundation for a flexible and extendable software package. These two principles guided our decision to use ParcPlace's VisualWorks application development environment and the OO language Smalltalk-80. VisualWorks allows the application developer to call C and C++ subroutines easily, which we need to do for computationally intensive work.

There are three components of the BIE that will be discussed in more detail: 1) the graphical programmer, which allows a user to graphically define and interact with a measurement model $y = H(x)$, 2) a geometric object manipulator, which allows a user to define and interact with geometric parameterizations of an object or collection of objects, and 3) an optimizer, which allows a user to coordinate the global optimization of a likelihood or posterior with respect to user-selected object parameters and get feedback on the progress of the optimization.

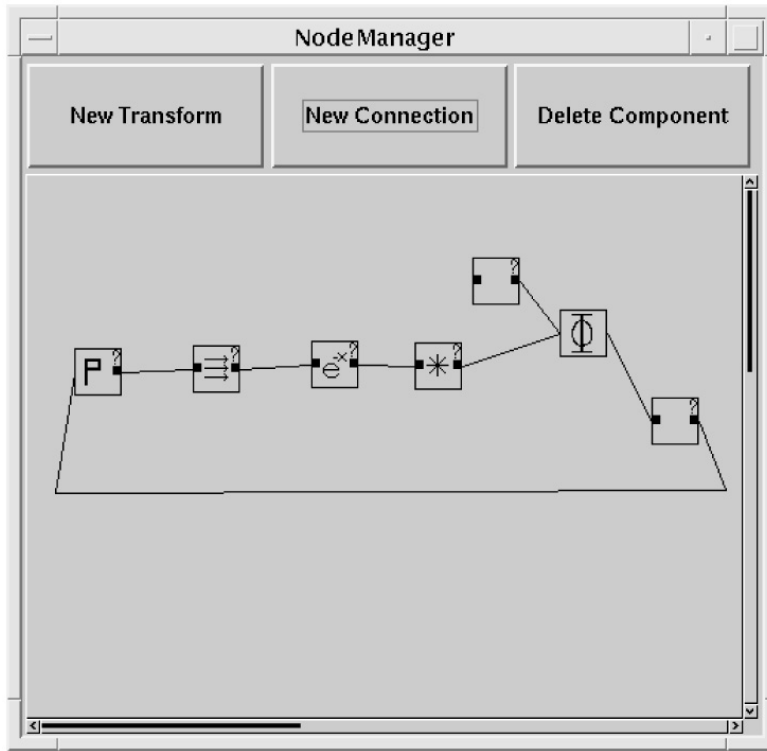


Figure 2. The canvas for the graphical programming tool showing, from left to right, a density image, projection transform, negative exponentiation, and convolution (to include detector blur). The icon labeled Φ computes the minus-log-likelihood from the inputs of the actual data (topmost icon) and the data calculated by the radiographic simulation model. The lower right icon is the Optimizer, which minimizes the output of the Φ box with respect to the image on the left.

2.2. The Measurement Model

Figure 2 shows a canvas on which a simple radiographic measurement system has been modelled using the graphical programming tool [1]. The tool allows a user to create, connect, delete, and reorganize icons that represent **Transform** objects. The lines between icons represent **Connector** objects that are capable of passing data forward and backward. **Parameter** objects have no input and a single output. Intermediate predicted data can be generated and viewed by telling any **Transform** to “generateOutput” and “displayOutput”, resp.. These messages are passed backward by the **Connectors** until a **Parameter** is reached and returns itself, the recursively transformed result eventually propagating backward to the initiating **Transform**. Similarly, the gradient of a **Parameter** with respect to a **Likelihood** object can be obtained by telling the **Parameter** to “generateAdjointOutput”. This message is passed forward by the **Connectors** until a **Likelihood** object is reached and returns the gradient of itself with respect to the predicted data, the transformed adjoint result eventually propagating back to the initiating **Parameter**. Thus the graphical programming tool can be used to construct a potentially complex measurement system H and likelihood ϕ acting on a parameterization x . Furthermore, the **Parameters**, **Connections** and **Transforms** are sophisticated enough so that ϕ and the gradient of ϕ with respect to any **Parameter** can be

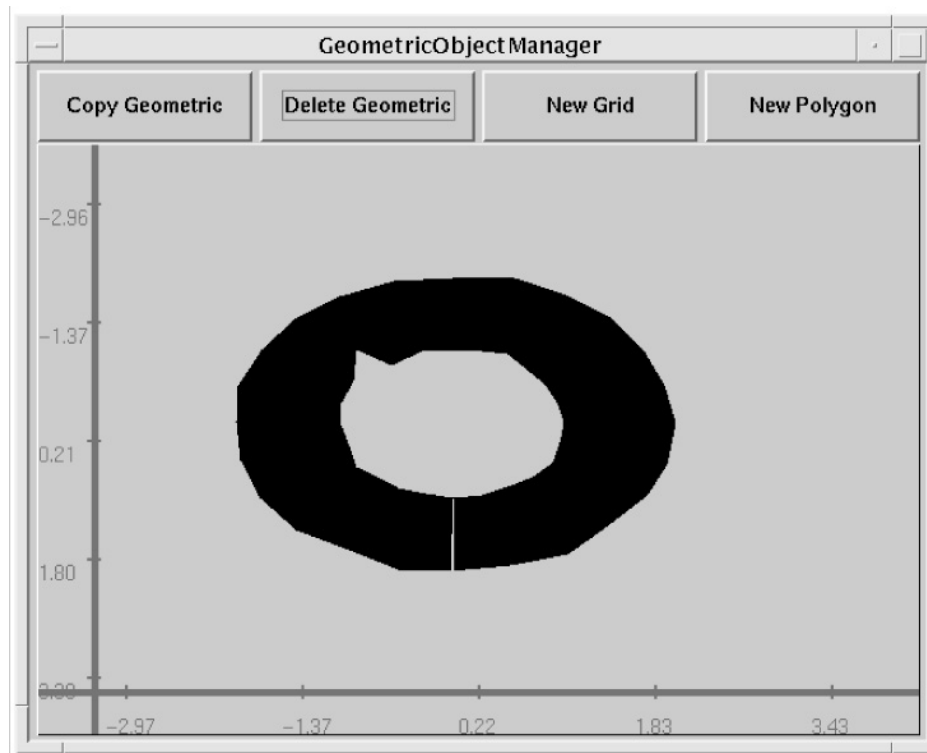


Figure 3. The canvas for the graphical programming tool showing, from left to right, a density image, projection transform, negative exponentiation, and convolution (to include detector blur). The icon labeled Φ computes the minus-log-likelihood from the inputs of the actual data (topmost icon) and the data calculated by the radiographic simulation model. The lower right icon is the Optimizer, which minimizes the output of the Φ box with respect to the image on the left.

computed.

2.3. The Geometric Object Manipulator

Figure 3 shows a BIE canvas on which a 2D geometric model of an experimental object has been constructed. Objects can be composed of many parts, each of which can have any one of several geometries, including `UniformGrid2D`, `Grid2D`, `Polygon2D`, and `PiecewiseBezier2D`. Object parts can be deleted, translated, resized, and modified appropriately in an interactive way. The values associated with `UniformGrid2D` and `Grid2D` can be seen in a grayscale image and manipulated with another tool. All geometric parts can be deformed using piecewise Bezier 2D warps. Contour plots of a `Grid2D` can be automatically refreshed as the `Grid2D` is interactively deformed. The geometric modeller is called when a `GeometricObjectParameter` is told to display itself (from the graphical programming tool).

2.4. The Optimizer

A user can tell a `Parameter` to be optimizable by connecting the `Parameter` icon to the `Optimizer` icon, of which there is usually only one on the graphical programming canvas. The code used

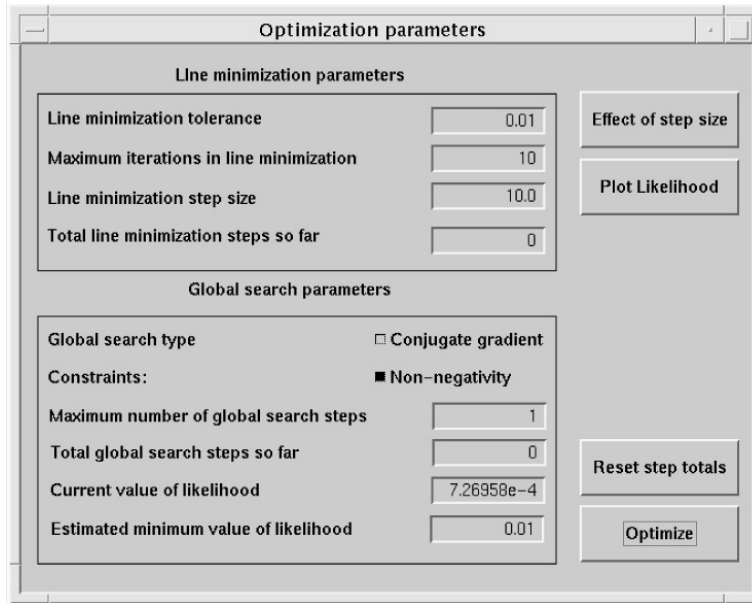


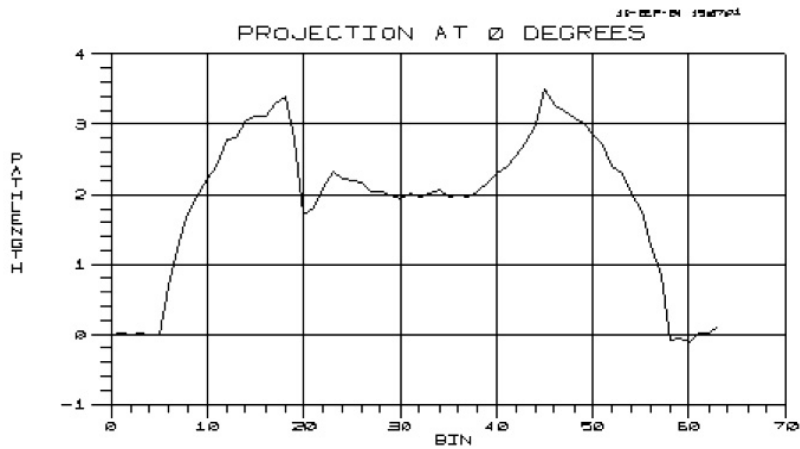
Figure 4. The user interface for the `Optimizer`.

by the `Optimizer` is abstract [2] in the sense that the `Parameters` it optimizes are not typecast, and details about the gradient calculation are not known by the `Optimizer` (this is the responsibility of the `Parameters`, as discussed above). `Parameters` also know how to multiply themselves by scalars, add themselves to a like-structured `Parameter`, find their inner-product with a like-structured `Parameter`, etc. The `Optimizer` uses the abstract vector-space operations, whose implementation is `Parameter`-dependent, to conduct constrained and unconstrained optimization using gradient descent, conjugate-gradient, and Powell's algorithms. Several line-search strategies have been investigated, including golden-section, polynomial fit, and a hybrid approach using both golden-section and quadratic fit. The user is presented with an interface that contains the important attributes of the optimization (see Fig. 4). The user can plot the effect of stepping in the current gradient direction on the optimizable `Parameters`. The user can also plot the likelihood as a function of step size along the gradient direction. Intermediate predicted data and the current state of the object model are always accessible during any optimization using the graphical programming tool and the geometric modeller.

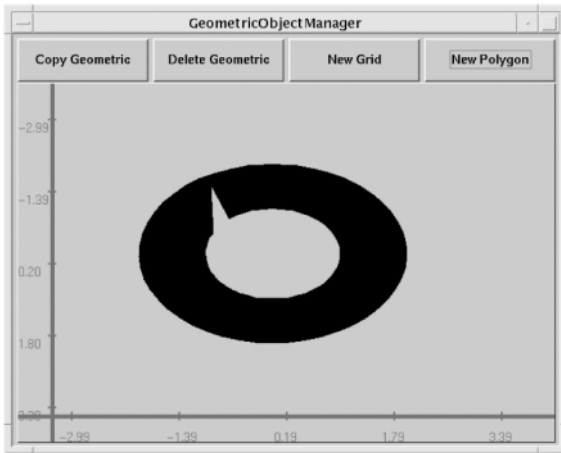
3. EXAMPLE

3.1. 2D Limited View Tomography


Figure 5(a) contains the geometry of the original attenuation profile, which is an annulus with an interior flaw. The inner radius of the annulus is 1 cm, the outer radius is 2 cm, and the attenuation value in the annulus is 1 cm^{-1} . The graphical programming tool is used to create the four noisy, simulated radiographic projections in Fig. 5(b) of the attenuation profile in Fig. 5(a). Each projection has 64 bins and a physical width of 5 cm. Random Gaussian noise with an rms value



(b)



(a)



(c)

Figure 5. (a) The original geometry: a 100-vertex polygonal approximation to an annulus with a large flaw on the interior. (b) The simulated vertical projection of the geometric object in Fig. 5(a). (c) The reconstructed geometry using a 50-vertex polygon with initial configuration equal to the unflawed annulus.

of 0.05 has been added to these simulated measurements, which represents a SNR of $\frac{3.4}{.05} = 64 = 18$ dB. The estimate in Fig. 5(c) is the result of optimizing the 50 vertices of a polygon with initial vertices lying on the unflawed annulus.

4. CONCLUSIONS AND FUTURE DIRECTIONS

The BIE is already a powerful tool for 2D geometry optimization, reliability investigation, and hypothesis testing; however, there are many unexplored avenues we intend to pursue. Two goals that we will be pursuing this year are extending the optimization capability of the BIE to more complex geometry and exploiting the BIE's interactivity to investigate global optimization strategies

that control the complexity of the object and measurement system models. We aim to eventually optimize 3D CAD geometry descriptions to fit limited radiographic data.

ACKNOWLEDGMENTS

Supported by the United States Department of Energy under contract number W-7405-ENG-36.

REFERENCES

1. G. S. Cunningham, K. M. Hanson, G. R. Jennings, Jr., and D. R. Wolf, "An object-oriented implementation of a graphical-programming system," in *Proc. SPIE* **2167**, M. H. Loew, ed., pp. 914-923, Newport Beach, CA, February 15-18, 1994.
2. G. S. Cunningham, K. M. Hanson, G. R. Jennings, Jr., and D. R. Wolf, "An object-oriented optimization system," in the *Proceedings of the International Conference on Image Processing*, Austin, TX, November 13-16, 1994.