

BAYESIAN MARS FOR UNCERTAINTY QUANTIFICATION IN STOCHASTIC TRANSPORT PROBLEMS

Hayes F. Stripling and Ryan G. McClarren

Dept. of Nuclear Engineering
Texas A&M University
3133 TAMU
College Station, TX 77843-3133
h.stripling@tamu.edu, rgm@tamu.edu

ABSTRACT

We present a method for estimating solutions to partial differential equations with uncertain parameters using a modification of the Bayesian Multivariate Adaptive Regression Splines (BMARS) emulator. The BMARS algorithm uses Markov chain Monte Carlo (MCMC) to construct a basis function composed of polynomial spline functions, for which derivatives and integrals are straightforward to compute. We use these calculations and a modification of the curve-fitting BMARS algorithm to search for a basis function (response surface) which, in combination with its derivatives/integrals, satisfies a governing differential equation and specified boundary condition. We further show that this fit can be improved by enforcing a conservation or other physics-based constraint.

Our results indicate that estimates to solutions of simple first order partial differential equations (without uncertainty) can be efficiently computed with very little regression error. We then extend the method to estimate uncertainties in the solution to a pure absorber transport problem in a medium with uncertain cross-section. We describe and compare two strategies for propagating the uncertain cross-section through the BMARS algorithm; the results from each method are in close comparison with analytic results. We discuss the scalability of the algorithm to parallel architectures and the applicability of the two strategies to larger problems with more degrees of uncertainty.

Key Words: Uncertainty Quantification, Stochastic Transport, Bayesian MARS, Stochastic solutions to PDEs

1. INTRODUCTION

One of the primary missions of scientific computing is the accurate solution of large-scale problems on the most advanced hardware platforms. To a large extent this mission has been successful as computational software now routinely takes advantage of tera- and petascale computers and in turn makes the solution of nonlinear problems with billions and trillions of unknowns routine. The further progress of computing capabilities will soon allow resolutions that are so refined that fundamentally uncertain features such as, for example, material grain boundaries are resolved by the spatial mesh. Given that the location of such features is practically unknowable in an as-built component, further refinement of the computational mesh will result in a very precise answer to the wrong problem. Treatment of uncertainty at this level will require new approaches—such as

methods for solving partial differential equations that treat all the parameters in the equations, including dependent and independent variables, as distributions.

At the same time there has been greater emphasis on using simulation to predict the behavior of complicated systems where experimental data is difficult or impossible to obtain for financial, safety, or political reasons. Examples of such situations are human space craft, nuclear reactor design, and stockpile stewardship. Using simulation to predict the behavior of such systems involves the synthesis of data obtained from small-scale or single physics experiments and simulation, as well as understanding how uncertainties in the simulation due to manufacturing tolerances, physical constants, and numerical error affect the prediction.

In this work we look at a new approach to solve particle transport problems where the cross-sections are uncertain, though, as we will demonstrate, it is applicable to a much wider range of problems. Generally, methods for propagating input uncertainties to output quantities of interest (QOIs) are classified as either intrusive or non-intrusive in regards to how they can be integrated with existing codes. Non-intrusive methods use existing codes to run series of calculations (often called run sets) and use the results of those calculations to infer distributions for the output QOIs. Examples of these methods are sampling techniques, collocation-based polynomial chaos [1–3], and reliability methods. Other methods are intrusive in that they require modifying existing codes or writing a new code to use. Common intrusive methods are stochastic finite element techniques [1] and automatic differentiation.

One non-intrusive method is surrogate-based sampling where a statistical regression model (such as Gaussian process regression [4] or multiple adaptive regression trees [5]) is built for an existing code. This model, or emulator, is then sampled based on the distribution of input parameters to get an estimate of the code's response to these input distributions. Multivariate Adaptive Regression Splines (MARS) has been shown to be effective as an emulator, especially for codes with high dimensional input or highly non-stationary response. The classical MARS algorithm [6] is a partition-based curve-fitting technique which attempts to emulate the mapping between a function's inputs and outputs as a summation of spline functions. The Bayesian extension of MARS (BMARS) [7] uses a Metropolis-Hastings algorithm [8] to construct an ensemble (or posterior distribution) of predictive MARS functions. The distribution of predictions generated from a BMARS evaluation can be used to propagate regression uncertainty and inherent uncertainty in the data to uncertainty in the estimated QOIs.

The effectiveness and construction of BMARS inspired the authors to investigate whether this technique could be used to estimate solutions to partial differential equations that have uncertain parameters. Such a method would solve the PDEs in a statistical sense and be able to handle inline any type of uncertain parameters. Insofar as this requires the design of an entirely new approach to solving PDEs as well as developing a new UQ method, our work is an über-intrusive UQ method and in the very early stages of development and exploration.

In this work we apply BMARS to solve simple slab-geometry particle transport problems where the material is a pure absorber of uncertain cross-section. In the following section we present BMARS and demonstrate how we apply it to transport problems. We then discuss how uncertainties in the material parameters are dealt with in Sec. 3 followed by numerical results in Sec. 4. In Sec. 5 we present our conclusions and discuss future work.

2. BAYESIAN MARS AS APPLIED TO TRANSPORT PROBLEMS

The MARS emulator constructs a mapping from model input to output using combinations of 1D polynomial spline functions. These spline functions are defined to be zero on part of the domain and a polynomial of some order on the remainder of the domain. The knot of the spline is the coordinate at which this definition changes, and the direction of the spline describes whether the non-zero portion of the spline is in the positive or negative direction from the knot. The BMARS basis function can be written as

$$B(\vec{x}) = \beta_0 + \sum_{k=1}^{\mathbf{K}} \beta_k \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}} \quad (1)$$

where \vec{x} is a vector of inputs, $t_{k,l}$ is the knot point in the l^{th} dimension of the k^{th} component, the function $(y)_+$ evaluates to y if $y > 0$, else it is 0, o is the polynomial degree in the l^{th} dimension of the k^{th} component, β_k is the coefficient of the k^{th} component, \mathbf{K} is the maximum number of components of the basis function, and \mathbf{I} is the maximum allowed number of interactions between the L dimensions of the input space.

BMARS generates an ensemble of predictive MARS functions which, in turn, provide a predictive distribution instead of a single prediction of the model response. Step zero of the algorithm generates a classical basis function in the form of Eq. (1). A Markovian process [9] proposes a change to the current model in the form of an addition, deletion, or modification of a spline. When a new spline is created, the algorithm randomly chooses its order, knot point, direction, and level of interaction. The algorithm iterates this random selection process and accepts/rejects proposal basis functions based on a likelihood calculation. This likelihood is a function of the candidate's fit to the training data and the number of splines in its basis function.

Fundamental to the present work is the simple formulation for the derivative of a BMARS basis function with respect to one dimension of \vec{x} . This derivative can also be expressed as a BMARS basis function:

$$\frac{dB(\vec{x})}{dx_n} = \sum_{k=1}^{\mathbf{K}} o_{k,n} \beta_k \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}^*} \quad (2)$$

where

$$o_{k,l}^* = \begin{cases} o_{k,l} - 1 & l = n \\ o_{k,l} & l \neq n \end{cases} .$$

Of course if a component of the basis function does not include a spline in the n^{th} dimension, then its derivative with respect to that dimension is zero. In the following paragraphs, we describe how we employ this simple formulation to estimate solutions to a simple transport problem.

We will consider the simple transport problem in 1D slab geometry with a pure absorber, no source, and isotropic incident flux on one boundary of the slab. The governing equation and

boundary condition are

$$\begin{aligned} \mu \frac{d\psi}{dz} &= -\sigma\psi(z, \mu) \\ \psi(0, \mu) &= 1 \\ z \in [0, Z], \mu &\in [0, 1] \end{aligned} \tag{3}$$

where z is the spatial coordinate, μ is the cosine of particle travel with respect to the z direction, σ is the interaction cross-section, and $\psi(z, \mu)$ is the particle angular flux. This problem treats only those particles traveling in the positive z direction, which in a pure absorber problem, is equivalent to treating all directions and specifying zero incident flux at $z = Z$. Problem 3 has analytic solution

$$\psi(z, \mu) = \exp\left(-\frac{\sigma z}{\mu}\right).$$

As merits of comparison, we will also be interested in the scalar flux and reaction rate in the slab, which are written respectively as

$$\phi(z) = \int_0^1 \psi(z, \mu) d\mu \tag{4}$$

$$R = \int_0^Z \sigma \phi(z) dz. \tag{5}$$

We attempt to construct an estimate for the solution to 3 using a modified BMARS algorithm. We seek a BMARS function which satisfies the following equation

$$\mu \frac{dB(\vec{x})}{dz} + \sigma B(\vec{x}) = \beta_0 \sigma + \sum_{k=1}^{\mathbf{K}} \beta_k \left[\sigma \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}^*} + \mu o_{k,nz} \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}^*} \right] = 0. \tag{6}$$

As the solution to the problem cannot be represented exactly by our polynomial splines, we transform the effort into a minimization problem. We first generate a Latin Hypercube sample of N coordinates in the independent input space (z and μ thus far). This sample is augmented with points at the boundary where we will enforce the boundary condition. The corresponding response vector is the zero vector of length N augmented by the ones vector corresponding to the boundary points. The LHS sample, response vector, and boundary information are analogous to the training inputs/response which are input to the curve-fitting BMARS algorithm. Finally, a matrix of discrete knot candidates in each input dimension is provided to the algorithm.

Similar to the curve-fitting BMARS algorithm, our program proceeds to construct a basis function by proposing, testing, and accepting/rejecting modifications to the current model. At a given step of the Markov chain Monte Carlo (MCMC) procedure, if a new spline is proposed, its derivative is also computed and each new spline (the proposed spline and its derivative) is

evaluated at the N sample coordinates. The result is two design matrices, each with N rows and K_j columns, where K_j is the number of splines in the basis function proposed by the j^{th} MCMC iteration. These matrices contain the evaluation of each spline at each of the N samples and correspond to the terms in the square brackets of Eq. (6).

Each MCMC iteration requires a solution for the K_j coefficients β , which we write as an over-determined linear system (we always assume that $N > K_j$). We solve for the coefficients β using standard constrained least squares theory. The coefficients are chosen to minimize the evaluation of the proposed basis function at the N sample points (Eq. (6)), but they are constrained such that the basis function also satisfies the boundary condition (enforcing the boundary condition also avoids finding the trivial solution, $\beta = 0$).

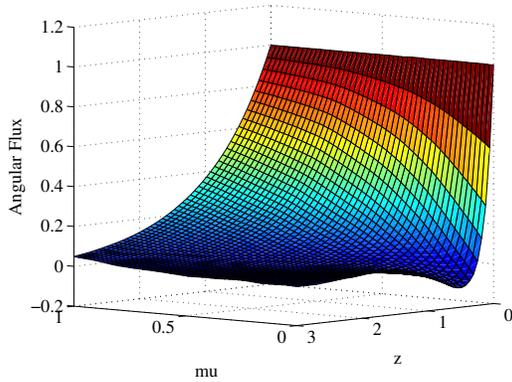
We employ a simple likelihood calculation, which we write as

$$L(B_j) = \left\| \mu \frac{dB_j(\vec{x})}{dz} + \sigma B_j(\vec{x}) \right\|_{L_2} f(K_j).$$

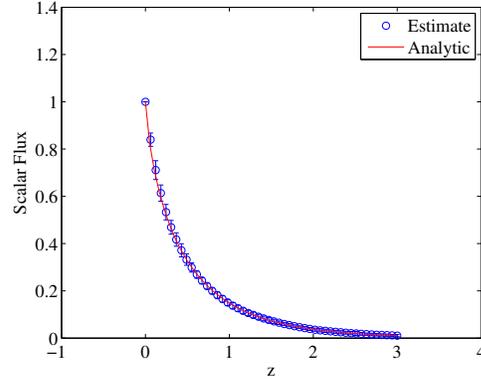
The first factor is simply the sum squares residual of the current estimate to the solution of our differential equation at the N LHS sample points. The second factor is related to the size of the current estimate. In some cases, we use a Poisson prior for K_j with mean λ to promote a basis function of certain size; in others we use a strictly decreasing gamma prior which will only grow larger if necessitated by the regression. The standard Metropolis acceptance algorithm [8] is used to accept or reject proposed basis functions.

Before moving to problems with uncertain cross section, we give results for estimates to the solution of the deterministic problem with $\sigma = 1$ and $Z = 3$. As a part of our algorithm, we compute and store BMARS integrals for scalar flux and reaction rate (Eqs. (4) – (5)). Integrals of BMARS functions are also BMARS functions, but are not as straightforward to compute as derivatives. We ran the algorithm for 5000 burn-in cycles and generated a posterior distribution of 5000 BMARS estimates for angular flux, scalar flux, and reaction rate.

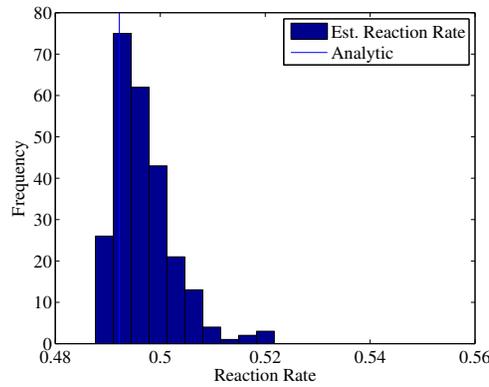
Figure 1 illustrates the result of sampling 250 random posterior functions. We make a number of observations. First, as there is no uncertainty in the problem itself (the exact solution to this problem is an exponential integral function), the distributions given in Figs 1(b) and 1(c) are strictly results of regression uncertainty; that is, uncertainty resulting from the search for the best fit to the governing differential equation. Second, we do see negative flux values in Fig. 1(a) near $\mu = 0$ for small z . The limit as μ approaches zero corresponds to particles moving perpendicular to the z axis, is a singular value in the analytic solution for the angular flux, and is generally difficult to characterize for most transport methods used today. We note that this is an unphysical result but, for the time being, do not give it special concern. Finally, reaction rate results greater than 0.5 in Fig. 1(c) violate a global particle conservation (the incident current is 0.5, and the reaction rate cannot exceed this value). These results prompted a study into the requirement of global conservation as a constraint on the least-squares solve for the BMARS coefficients β . We expand upon this study below.



(a) Angular flux in $\{z, \mu\}$ space



(b) Scalar flux (90% confidence intervals) compared to the analytic solution.



(c) Reaction rate distribution compared to analytic

Figure 1: Results from a simple BMARS estimate for the solution to problem 3.

3. TREATING UNCERTAIN CROSS-SECTIONS

The method described above will be most useful when applied to differential equations with uncertain parameters. As an illustration, we introduce a modification to problem 3:

$$\begin{aligned}
 \mu \frac{d\psi}{dz} &= -\sigma\psi(z, \mu) \\
 \psi(0, \mu) &= 1 \\
 \sigma &\sim \Gamma(16, 0.05) \\
 z &\in [0, Z], \mu \in [0, 1].
 \end{aligned}
 \tag{7}$$

The cross-section is now assumed to be uncertain, but we estimate that it follows a Gamma distribution with mean of 0.8 and variance of 0.04 (a Gamma-distributed random variable is strictly non-negative). Note that σ is still assumed to be constant in the slab, not spatially dependent. The end goal is to produce a posterior distribution of predictive functions that contain

information about the cross-section uncertainty and propagate this uncertainty to calculations of the scalar flux and reaction rate.

We present two methods for treating the uncertain cross section. The first method (method A) involves the addition of σ as a third input dimension in the BMARS model, making it a free parameter in the same sense that z and μ are free parameters. The N input LHS points now include samples from $\sigma \in [0.5, 2.0]$ (this range covers more than 99% of the Gamma pdf), the algorithm is allowed to choose knots and generate splines in the σ dimension, and the resulting BMARS models for angular flux, scalar flux, and reaction rate are a function of μ, z , and σ . We then produce realizations of these models (*i.e.* generate a distribution of the outputs) by sampling our uncertain Gamma input distribution for σ .

In the second method (method B), we do not include σ as a free parameter, but instead sample from its uncertain distribution during the MCMC process. We first burn-in a model with constant σ equal to the mean of its uncertain distribution. We then sample a value of σ from the Gamma distribution and run the MCMC process for N_t iterations to allow it to adapt to the perturbation. We save the last N_s of these iterations before re-sampling a new value of σ and repeating. Appropriate values of N_t and N_s are still under investigation, but certainly are problem dependent and could be determined heuristically or by converging a “goodness of fit” parameter after each perturbation before saving and moving on to the next realization.

The posterior distribution resulting from method B is made of blocks (each of size N_s) of realizations that correspond to a sampled value of σ . The mean prediction of the N_s realizations from each block can be used as an estimate for the BMARS prediction corresponding to that value of σ . Care must be taken to sample a sufficient number of σ s to fully explore its distribution.

The fundamental difference between the methods is that method B samples the uncertain parameters during the model construction while method A samples uncertain parameters after model construction. Also, method A requires an extra dimension in the model for each uncertain parameter (*i.e.*, if we had both an uncertain σ and an uncertain Z , method A requires two additional dimensions to place knot points). In the case where there is a large number of uncertain parameters in the problem, this could present a challenge as the number of knot points and dimensions for sampling would need to be large. In method B the values of N_s and N_t will, in all likelihood, need to be large when there are many uncertain parameters. Method B, however, is very amenable to parallel computing as the samples of the uncertain parameters can be processed independently after the initial burn-in.

4. RESULTS

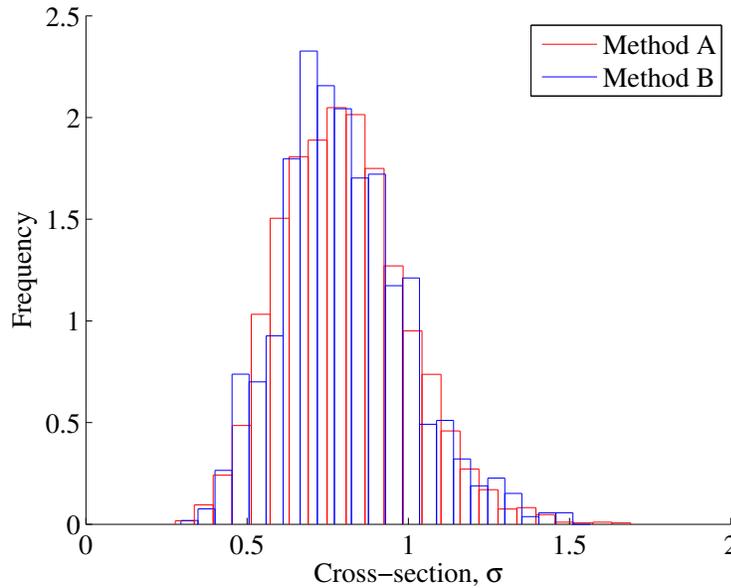


Figure 2: Samples from the uncertain distribution of σ for methods A and B.

In the interest of comparison, the results from methods A and B will be presented side-by-side. The three-dimensional model for method A was constructed with a 35,000 iteration burn in, and 5,000 posterior realizations were saved. We then sampled 1,000 values of σ to produce the statistics given in the following plots. In method B, we sampled 1000 values of σ , with $N_t = 100$ and $N_s = 10$. The results given for method B are related to the mean of the 10 saved realizations for each value of σ . Figure 2 shows the distribution of the sampled σ s for the two methods.

The statistics of interest are the distribution of the spatially dependent scalar flux in the slab as well as the distribution of the reaction rate. The distributions labeled “analytic” are produced by computing the analytic scalar flux and reaction rate for a large number of samples of σ from its uncertain distribution. The distributions labeled BMARS are generated by methods A and B as described in the previous section.

Figure 3 shows the spatially dependent scalar flux distributions for methods A and B compared to the analytic distribution. The mean and 90% confidence interval of each method is in close agreement with the analytic calculations (method B tends to slightly overpredict). One interesting feature is the behavior of the 90% confidence interval around the left boundary of the slab, where the boundary condition is enforced at a discrete number of points. For methods A and B, a boundary coordinate is written as $\{z = 0, \mu, \sigma\}$ and $\{z = 0, \mu\}$, respectively. The additional input dimension in method A requires the constraint of a function space, which we approximate by enforcing the boundary condition at a number of sampled points on the $z = 0$ boundary. Incomplete constraint, however, results in more predictive variance at the boundary.

For a more quantitative comparison, Fig. 4 shows the error in the mean prediction of each

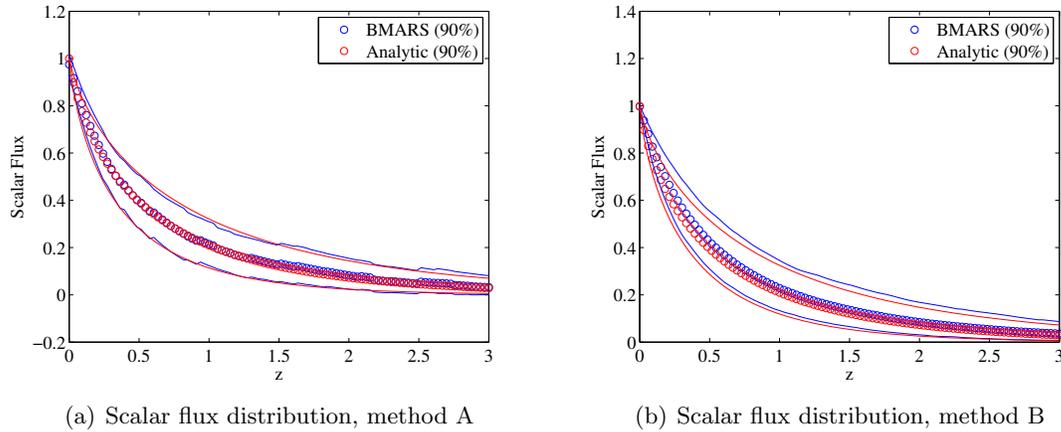


Figure 3: The BMARS and analytic distributions for scalar flux using the two sampling methods.

method. Both methods tend to overpredict the scalar flux, with the root-mean-squared predictive error of methods A and B being 0.0154 and 0.0277 respectively.

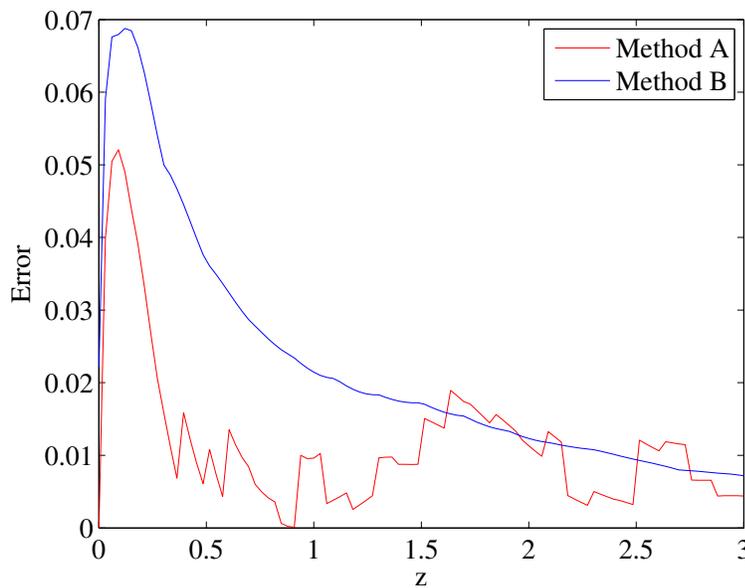


Figure 4: Error in the mean prediction of methods A and B.

Figure 5 shows the distribution of reaction rates for the two methods compared to the analytic distribution. The top row of figures gives the full range of computed reaction rates, and row two of the figures zooms-in on the analytic range of reaction rates to resolve the features of the BMARS distributions.

In regards to the reaction rate, the means of the distributions produced by methods A and B

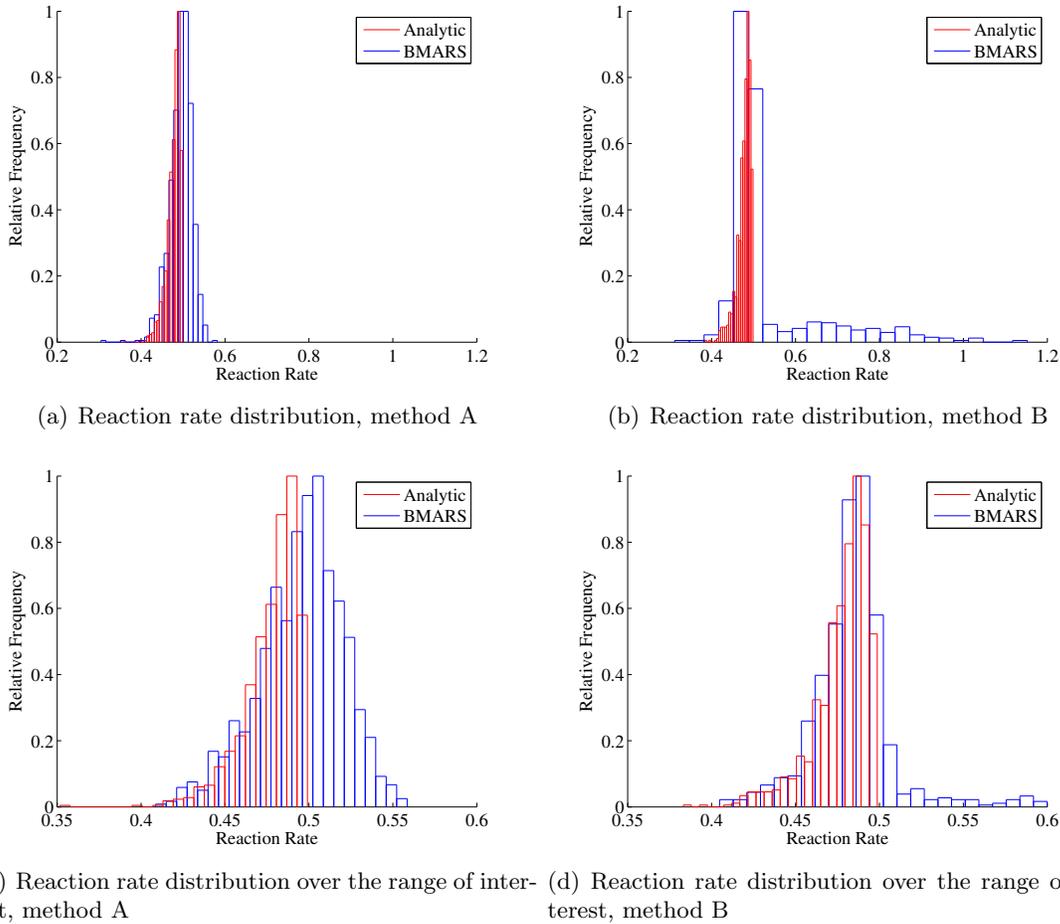


Figure 5: The BMARS and analytic distributions for reaction rate using the two sampling methods.

are 0.495 and 0.532 respectively, while the analytic mean is 0.478. Both methods overestimate the reaction rate, which is expected as each method tended to overpredict the scalar flux. The variance in the distribution corresponding to method A is smaller than that of method B, as method B has a large tail to the right of its mode. The majority of the distribution, however, more closely resembles the analytic distribution than does the distribution of method A. We believe method B has more inherent variability, as the model is changing during the basis construction. Future work should include more thorough treatment of model convergence as the uncertainties in the problem are sampled. For example, the current algorithm assumes that the fixed values of N_t and N_s are appropriate to accommodate changes in the governing equation.

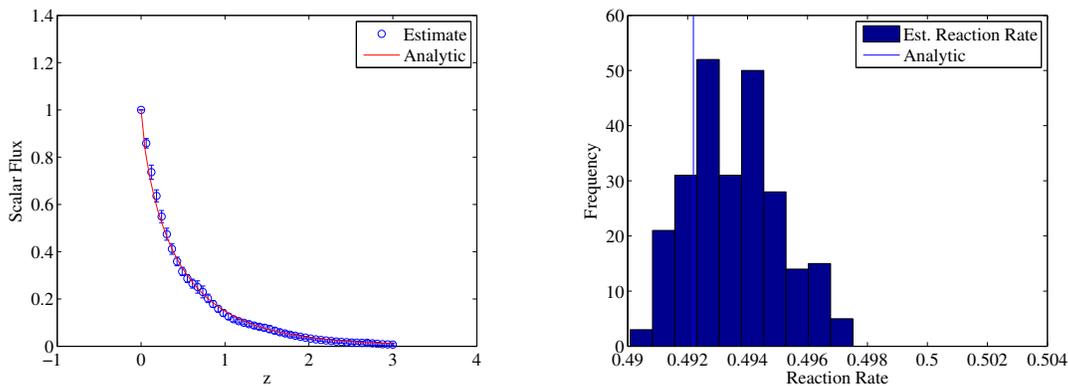
As mentioned above and seen again in Fig. 5, global particle conservation is not met when our model predicts reaction rates greater than the magnitude of the incident current, 0.5. The ability to enforce and show conservation would strengthen the legitimacy of our model and possibly improve its accuracy by injecting a physical constraint into the UQ problem. In terms of transport quantities, the conservation statement for problems (3) and (7), which simply says that the sum of the exiting current and reaction rate must equal the incident current, is written as

$$\int_0^1 d\mu \mu \psi(z=0, \mu) - \int_0^Z \sigma(z) dz \int_0^1 d\mu \psi(z, \mu) - \int_0^1 d\mu \mu \psi(z=Z, \mu) = 0.$$

Written in terms of BMARS quantities, the conservation statement becomes an additional equation or constraint in the over-determined least-squares solve for the β s:

$$\sum_{k=1}^{\mathbf{K}} \beta_k \left\{ \int_0^1 d\mu \mu \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}} \Big|_{z=0} - \int_0^Z \sigma(z) dz \int_0^1 d\mu \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}} - \int_0^1 d\mu \mu \prod_{l=0}^{\mathbf{I}} (x_l - t_{k,l})_+^{o_{k,l}} \Big|_{z=Z} \right\} = 0.$$

Figure 6 illustrates the results of adding this constraint and re-computing an estimate to the deterministic problem (problem (3)). We see roughly the same regression noise in the estimate for the scalar flux, but a *much* tighter distribution of predicted reaction rates about the analytic value and no unphysical values larger than the incident current. Again, the predictive distributions in these plots are representative of the regression error inherent to the BMARS model.

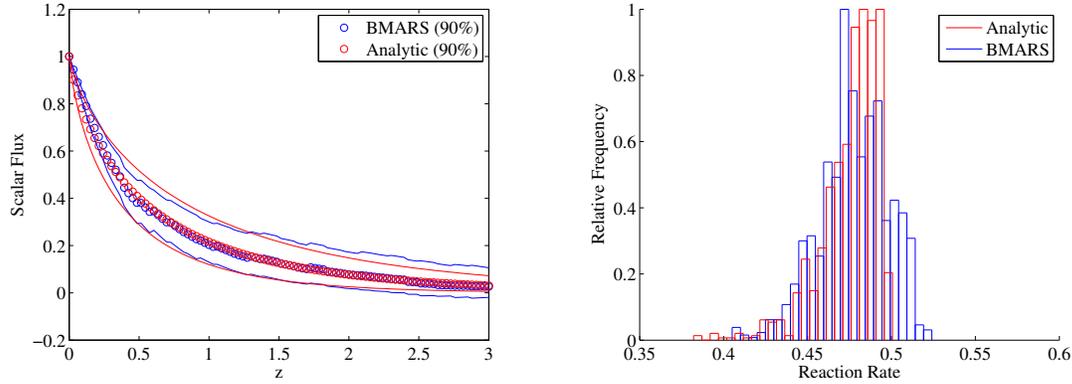


(a) Predicted scalar flux with conservation constraint. (b) Predicted reaction rate distribution with conservation constraint.

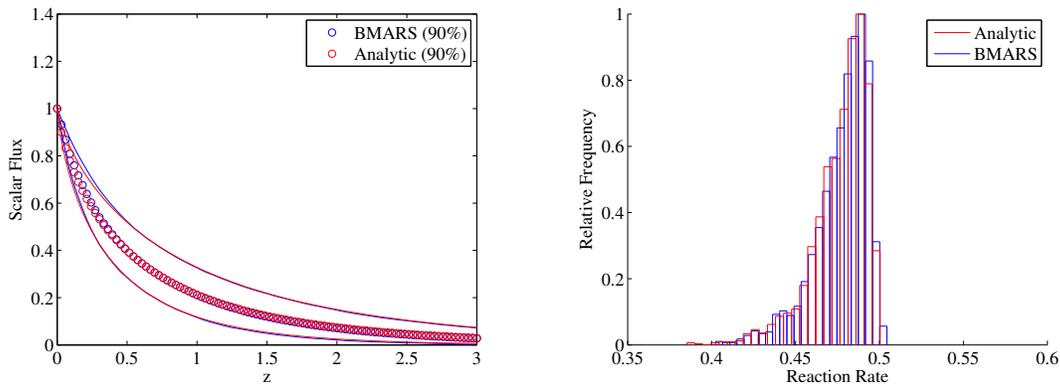
Figure 6: Results for the deterministic problem with the global conservation constraint on the model.

Next we applied the conservation constraint in methods A and B for the problem with uncertain cross-section (problem (7)). Implementing the constraint in method A again requires the constraint of a function space, which we at present are unable to do. Instead (as in the case of enforcing the boundary conditions in method A) we choose a discrete set of σ values at which to evaluate and enforce the conservation statement. The implementation in method B is nearly as simple as in the deterministic case, with only the requirement that splines generated in the burn-in are reevaluated each time a new random value of σ is drawn.

Figure 7 illustrates the scalar flux and reaction rate distributions predicted by the conservative implementations of methods A and B.



(a) Scalar flux distribution – conservative method A (b) Rx rate distribution – conservative method A



(c) Scalar flux distribution – conservative method B (d) Rx rate distribution – conservative method B

Figure 7: Predicted scalar flux and reaction rate distributions given by conservative implementation of methods A and B.

Compared to the results without conservation, the results given in Fig. 7 are more accurate in predicting the mean and variance of the scalar flux and reaction rate. Method B is now more accurate than method A and has only a single realization with a reaction rate greater than the incident current (which likely to be a result of numerical treatment to avoid a singular system in the constrained least-squared solve). In general, we see more noise and unphysical reaction rates in method A, which again are a result of its higher-dimensionality.

5. SUMMARY, CONCLUSIONS, AND FUTURE WORK

We have outlined and provided results from a method for estimating solutions to differential equations with uncertain parameters using a Bayesian MARS framework. We have shown that the method provides accurate answers that are within small regression error for a simple transport equation (without uncertainty) and extended the idea via two different strategies for handling uncertainty in the material cross-section. Therefore one conclusion of this work is that the BMARS approach to approximating solutions to simple differential equations may be useful when analytic solutions are difficult to compute. We also showed that the application of

a physical constraint, such as particle conservation, to the UQ problem improved the accuracy and legitimacy of the solution method.

Constructing BMARS solutions to PDEs for parallel computing is obviously an important issue for future work. Each MCMC iteration in the construction of the solution requires the solution of a dense, linear system of a size proportional to the number of knot points. Therefore, there is a practical limit to the size of problem we can solve with the standard BMARS modality. We have begun to devise parallel strategies for constructing BMARS solutions that involve a treed structure. In this approach the domain is decomposed and distributed to different processors. On each processor a BMARS solution is constructed and these BMARS solutions are connected to the entire domain by a global BMARS model where each processor is a knot point. Other strategies are possible and are being considered by the authors in ongoing work.

Our two strategies for handling a single uncertain parameter produced similar results for this problem. The authors believe that future studies involving a larger number of uncertainties will require an approach similar to method B. This method allows for sampling of the uncertainties *within* the model construction instead of posterior sampling. Further, method A requires that an uncertain parameter be treated as an independent input in the model construction, which would not be possible for many common forms of uncertainty. We believe that sampling from multiple sources of uncertainty during the model construction is a scalable and efficient method for generating posterior realizations of model predictions.

Ongoing and future work using this approach are numerous. The authors have new results from transport problems with scattering, which requires the solution to an integro-differential equation, as well as transport problems with fixed or spatially dependent source. The results are too preliminary for inclusion in this communication but are comparable to discrete ordinate calculations and may be included in the oral presentation at conference. The authors also hope to explore parallelization strategies for this method, especially in implementations similar to method B as described above.

ACKNOWLEDGEMENTS

Funding for the first author was provided by the Department of Energy Computational Science Fellowship under grant number DE-FG02-97ER25308. The work of the second author was supported by the DOE NNSA/ASC under the Predictive Science Academic Alliance Program by grant number DEFC52-08NA28616.

REFERENCES

1. Erin D. Fichtl, Anil K. Prinja, and James S. Warsa. Stochastic methods for uncertainty quantification in radiation transport. In *Proceedings of the International Conference on Mathematics, Computational Methods and Reactor Physics*, Saratoga Springs, NY, May 2009. American Nuclear Publication.
2. Erin D. Fichtl, Anil K. Prinja, and James S. Warsa. Uncertainty quantification in radiation transport using the stochastic collocation method. *Trans. Am. Nucl. Soc.*, 101:1106–1107, 2009.

3. Erin D. Fichtl, Anil K. Prinja, and James S. Warsa. The stochastic collocation method for radiation transport in random media. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 112(4):646–659, March 2011.
4. A. O’Hagan. Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety*, 91:1290–1300, 2006.
5. Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 28, 2000.
6. Jerome H. Friedman. Multivariate adaptive regression splines (with discussion). *The Annals of Statistics*, 19(1):1–141, 1991.
7. David G. T. Denison, Bani K. Mallick, and Adrian F. M. Smith. Bayesian MARS. *Statistics and Computing*, 8(4):337–346, 1998.
8. S. Chib and E. Greenberg. Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335, 1995.
9. L. Tierney. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4):1701–1728, 1994.