Bayesian Analysis of Inertial Confinement Fusion Experiments at the National Ignition Facility

J. A. Gaffney

June 22, 2012

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Bayesian Analysis of Inertial Confinement Fusion Experiments at the National Ignition Facility

J.A. Gaffney[a,*], D. Clark[a]

[a]*Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, CA 94550*

Abstract

We develop a Bayesian inference method that allows the efficient determination of several interesting parameters from complicated high-energy-density experiments performed on the National Ignition Facility (NIF). The model is based on an exploration of phase space using the hydrodynamic code HYDRA. A linear model is used to describe the effect of nuisance parameters on the analysis, allowing an analytic likelihood to be derived that can be determined from a small number of HYDRA runs and then used in existing advanced statistical analysis methods. This approach is applied to a recent experiment in order to determine the carbon opacity and X-ray drive; it is found that the inclusion of prior expert knowledge and fluctuations in capsule dimensions and chemical composition significantly improve the agreement between experiment and theoretical opacity calculations. A parameterisation of HYDRA results is used to test the application of both Markov chain Monte Carlo (MCMC) and genetic algorithm (GA) techniques to explore the posterior. These approaches have distinct advantages and we show that both can allow the efficient analysis of high energy density experiments.

Keywords:

1. Introduction

High energy density experiments conducted on large laser facilities are often highly complex in every way. The physics of the target's interaction with the laser, its evolution with time, the relevant material properties at the conditions of interest, and the relation between target properties and observable quantites all require significant effort to describe, and all involve approximations. This is particularly true of so-called 'integrated' experiments, such as the inertial confinement fusion (ICF) effort underway at Lawrence Livermore National Laboratory [1, 2], in which extremely complex targets probe all aspects of the ICF approach to nuclear fusion.

In general the only way that such complex systems can be accurately understood is through computer simulation. These simulations take a range of inputs desribing the target, laser system, and input physics and provide a description of the target evolution. This can then be mapped onto the output of experimental diagnostics. The approach represents a nonlinear, 'black box' transformation from the uncertain input physics to the measured quantites; the inversion of this to provide information about difficult physics can be a very difficult statistical problem. In this paper we discuss the application of Bayesian statistics [3] to the solution of this problem, and in particular to the inference of interesting physical parameters from an ICF experiment performed at the National Ignition Facility (NIF).

The Bayesian approach is a well developed statistical method that can be formulated to give powerful results for the inference of interesting quantities [4–6], or alternately to provide a rigorous framework for the design of new experimental methods [7–9]. Bayesian theory includes the prior knowledge of the system explicitly and so allows the evolution of interesting quantities to be charted as data accumulates from various sources. The numerical application of the theory is usually through Markov chain monte carlo techniques [10]; in the below we develop a method of incorporating the advantages of the Bayesian theory into genetic algorithm approaches which can be parrallelised very easily. These numerical methods have been applied in almost all areas of science for data analysis, including fusion studies, high energy density physics experiments [11], and hydrodynamic simulations [4, 12]. The example chosen here includes aspects of all of these types of experiment.

In the following section the Bayesian approach is set out. Particular attention is paid to the interpretation of Bayes' theorem as an information processing rule, since this can be very useful in understanding the evolution of physical models. The relationship between Bayesian methods and the common χ^2 methodology is discussed, which allows a generalised χ^2 to be defined that takes into account the effect of nuisance parameters in an experiment, and of prior knowledge of the system. This is applied to a NIF 'convergent ablator' experiment, and used to show the variation in the expected carbon opacity as these Bayesian factors are introduced. Finally, the MCMC and GA approaches are applied to the same analysis and discussed.

2. The Bayesian approach

In the Bayesian approach real numbers are used to describe the level of belief that a statement is true. A set of intuitive axioms govern these numbers, which coincide with the axioms of the usual probability theory. As such, Bayesian statistics can be developed as a re-intepretation of frequentist statistics; rather than the usual understanding that the probability $P(X = x) = \text{Lim}_{N \to \infty} \frac{N(X=x)}{N}$, $P(X = x)$ becomes the belief that the statement $X = x$ is true. This belief depends on some background information $\mathbb{I}$, a dependance that is expressed using the conditional probability $P(X = x|\mathbb{I})$.

An experiment has the effect of updating the background information with some observation $D = d$, so that after the experiment the belief becomes $P(X = x|D = d, \mathbb{I})$. The analysis of this updated probability forms the basis of Bayesian analysis. The axioms of the theory allow the important Bayes' theorem to be derived,

$$P(x|d, \mathbb{I}) \propto P(d|x, \mathbb{I})P(x|\mathbb{I}), \tag{1}$$

which relates the probability after the observation of D (the *posterior*) to the probability before (the *prior*) through the probability of observing the data when the value of X is assumed, $P(d|x, \mathbb{I})$ (the *likelihood*). The terms in equation (1) can be viewed as functions of any of their arguments; the study of $P(x|d, \mathbb{I})$ as a function of x allows the value of x to be inferred from data, whereas treating $P(x|d, \mathbb{I})$ as a function of d allows the significance of particular observations to be explored. The latter of these allows the design of experiments by ensuring that an experiment collects the most significant data possible.

Bayes theorem can be viewed as rule for describing the flow of information in an experiment [13]. The *information* associated with an observation of the random variable Z is

$$I(Z = z) = -\text{ln}P(Z = z) , \tag{2}$$

where the above is measured in 'nats'. The more usual unit of 'bits' is found using a base 2 logarithm. The information associated with an unlikely observation is greater than that with an unsurprising result. Bayes' theorem becomes

$$I(x|d, \mathbb{I}) = I(d|x, \mathbb{I}) + I(x|\mathbb{I}) + \text{constant} , \tag{3}$$

and so so the effect of an experiment is to add the new and prior informations. The information defined in equation (2) is evaluated for a given observation result z; when this is not known (as in problems of experimental design) it is necessary to consider the expectation value of the information, the *information entropy*

$$H[P] = - \int P(z)\text{ln}P(z)dz , \tag{4}$$

which is now a property of the distribution P only. The *cross entropy* between two distributions P and Q,

$$H[P\|Q] = - \int P(x)\text{ln}Q(x)dx , \tag{5}$$

measures the common information in the two distributions. This is particularly useful since it is conserved in Bayes' theorem [13]. Using the above definition equation (1) can be written as

$$H\left[P(x|d, \mathbb{I}) \left\| \frac{P(x|d, \mathbb{I})}{P(x|\mathbb{I})} \right. \right] = H\left[P(x|d, \mathbb{I})\|P(d|x, \mathbb{I})\right] ,$$

or

$$H_{KL}\left[P(x|d, \mathbb{I})\|P(x|\mathbb{I})\right] = H\left[P(x|d, \mathbb{I})\|P(d|x, \mathbb{I})\right] , \tag{6}$$

2

so that the increase in cross entropy between the posterior and prior (the *Kullback-Leibler distance*) is equal to the cross entropy between posterior and likelihood. Experimental design problems often maximise H_{KL} to maximise the effectiveness of an experiment. Using the above definitions it has been shown that Bayes' theorem is the only information processing rule that exactly conserves information cross entropy [13].

The above definitions are all useful forms of the same rule; the information that a given experiment provides is described by the likelihood function $P(d|x, \mathbb{I})$. Data analysis is then the analysis of the likelihood function for a given experiment and model, and standard models can be understood as special cases of the likelihood and prior. For example, it is common to use a normally distributed likelihood

$$P(d|x, \mathbb{I}) \propto e^{-\frac{1}{2}\left(\frac{d - \hat{d}(x)}{\sigma}\right)^2},$$

where $\hat{d}(x)$ is a model for the data given the parameter x; in this case the information associated with an observation is the usual χ^2 function [3]. The χ^2 minimisation method is then equivalent to finding the value of the parameter x that minimises the information provided by the experiment (that is, the parameter for which the observation is the least surprising). For a constant prior probability $P(x|\mathbb{I})$ (often used to describe complete ignorance) this is in turn equivalent to finding the value of x that has the largest posterior probability. Hence the χ^2 minimisation method can be seen to contain two approximations; the neglect of the structure of the posterior, and the neglect of prior information.

The prior information plays an important role in Bayesian analysis. In general it contains all information that is available before the experiment; expert opinion, the results from previous experiments, etc. Its inclusion allows otherwise very difficult problems to be tackled, and allows the belief in a model or parameter to evolve over time by using the posterior from one experiment as the prior of another. As such data analysis can include the results from entire campaigns of experiments common in modern science, and these experiments need not have the same design or be performed by the same people. All that is required is an expression for the likelihood function for a given experimental design.

3. Calculation of the likelihood function from large scale hydrodynamic simulations

In this work we aim to investigate the Bayesian analysis of inertial confinement fusion experiments performed at the National Ignition Facility. The usual analysis procedure is to use radiation-hydrodynamic simulations to provide post-shot analysis; measured quantities like laser power are used as inputs for simulations, the outputs of which are compared to other measured quantities. This comparison is used to gain information about the physics used by simulations, and other unmeasured properties of the experiment.

We consider NIF shot N110625, a 'convergent ablator' experiment designed to mimick the implosion dynamics of a full ignition shot. This shot is of interest since large-scale post shot simulations have been performed in the past, in which a large number of target parameters were varied [15]. This dataset will serve to benchmark the smaller set of simulations used in this study.

3.1. HYDRA Simulations

The eventual aim is that advanced numerical methods such as Markov-Chain Monte Carlo (MCMC) routines, or Genetic Algorithms (GA), are used to perform efficient analysis of NIF data. To acheive this a robust platform for running hydrodynamic simulations is required, and the feasibility of such large-scale simulations must be investigated. To that end a set of fortran wrapper routines have been developed that allow the radiation-hydrodynamic code HYDRA to run as a subroutine to a generic MCMC or GA code. These routines take a set of values of various input parameters (opacity, equation of state and drive modifiers, etc.) and return various measurable characteristics of the implosion. Modification of input physics databases, and extraction of implosion dynamics, is performed using existing scripts to ensure that returned values are consistent with those commonly used in NIF data analysis. Calculations are done in parallel, allowing an analysis code to consider several tens of thousands of HYDRA simulations in one job allocation.

In this work a grid of 1024 simulations of a convergent ablator implosion has been generated, in which a constant multiplier is applied to the absorption opacity of carbon and to the X-ray drive. A representation of the results is shown in figure 1, in which the velocity of the centre of mass of the fuel is plotted as a function of carbon opacity multiplier (averaging over all other parameters). Black points show the results of the earlier parameter scan, with the blue line and shading representing a moving average of the mean and standard deviation of those results [15]. The mean and standard deviation of the results of this work are shown in red. The solid line represents the nominal (unmodified drive) case, and the broken line shows the drive-averaged case. The large difference between these lines demonstrates
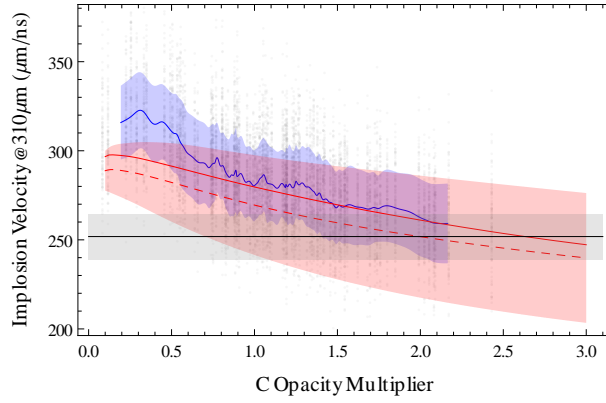
Figure 1: The implosion velocity as a function of carbon opacity multiplier for HYDRA simulations of N110625. Blue and red regions show one standard deviation about the mean of previous design simulations and this work respectively. Black shading indicates the experimental error bar.

the difficulty in increasing implosion velocity in these targets. Finally, the black region shows the experimental value [16].

The data in figure 1 agree well for the nominal case. The large number of parameters in the original data results in statistical noise in these results, the non-overlapping set of modifications made to simulations (original simulations focused more on changes to the capsule design than input physics), and the differences between the sampling of parameter space mean that the averages do not agree. The good agreement between the current work's nominal case and the previous results demonstrates that the HYDRA simulations used here are reasonable. It should be noted, however, that the simple modifications to the drive used here do not reflect the current understanding of the X-ray drive in ICF experiments, rather they were chosen as the drive multiplier is likely to correlate with the opacity multiplier.

3.2. Calculating the Likelyhood

We will use three measured results in the calculation of the likelihood; the implosion velocity when the centre of mass of the fuel reaches a radius of $310\mu m$, the time at which this radius is reached, and the fraction of ablator mass still remaining at this time [16]. Experimental data for these are available, and they can be easily extracted from HYDRA simulations. There are no simulated points that are able to match all three experimental values, making a statistical approach essential.

As discussed in the previous section, a simple model for the likelihood is an uncorrelated multivariate normal distribution. The information associated with the likelihood in this case, as a function of the two parameters we wish to infer, is plotted in figure 2. The most likely values of the drive and opacity multipliers occur at the minimum point of this surface, when the data are the least surprising. If each of the 3 experimental points are considered individually, each one defines a range of values of opacity and drive multipliers that best match the experimental result. The minimum information point in figure 2 results from the superposition of the three data points, giving rise to a broad range of (approximately) equally likely values of the simulation input parameters.

The assumption of a normal distribution for the calculation of the likelihood is based on assumptions about the nature of the *experimental* error. Simulations are treated as deterministic. In practice this is not the case; the values of input parameters will have a probability distribution as well (reflecting the accuracy with which they are calculated, for example) and the resulting distribution of code outputs will change the shape of the likelihood function. Proper inclusion of all sources of error rapidly increase the parameter space for simulations and so it is advantageous to use an approximate method for some parameters. For this we use a normal linear model [8], which we develop below.

Consider the likelihood function for a simulation with two inputs as before, but where one is a 'nuisance' parameter. If the nuisance parameter is normally distributed, and the response of the simulation to variations in the parameter is linear, then the likelihood function is a normal distribution with a modified width (see below). This linear response method can be used to include the effect of a large number of nuisance parameters with a greatly reduced number of simulations. The potential accuracy of this approach is improved by the fact that input physics models, target fabrication, etc. are highly developed for ICF applications making the expected variations in nuisance parameter values small.
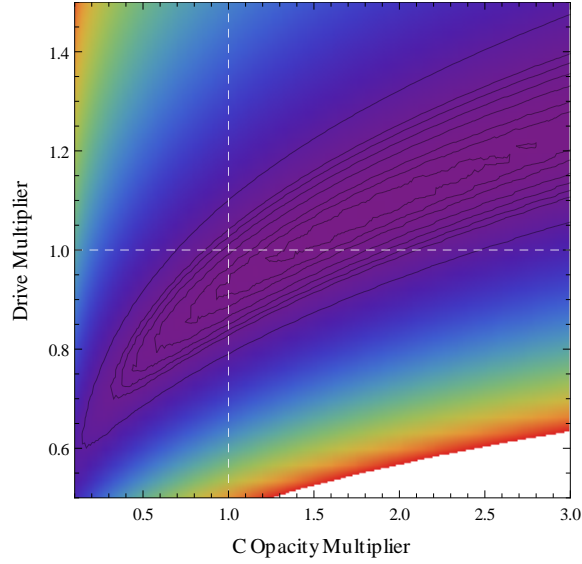
4

Figure 2: The likelihood as a function of simulation parameters, calculated using a multivariate normal distibution. Contour lines are logarithmic and centered on the minimum information parameters

To develop the linear response model, consider a normal approximation in which the likelihood is given by the multivariate distribution with covariance matrix Λ_d,

$$P(d|x, \mathbb{I}) = \frac{e^{-(d-\hat{d}(x))^T \Lambda_d^{-1}(d-\hat{d}(x))}}{\sqrt{|\Lambda_d|(2\pi)^{n_d}}} \ .$$

Splitting the simulation input vector x into the set of interesting parameters x' and nuisance parameters y, introducing a second multivariate normal distribution (of covariance Λ_y) to describe the variations in nuisance parameters, and treating the response of a simulation as linear,

$$\hat{d}(x) = \hat{d}(x', y) \simeq \hat{d}(x, \bar{y}) + A(y - \bar{y}) \ , \tag{7}$$

the likelihood function can be found analytically. Using $P(d|x') = \int P(d|x', y)P(y)dy$, the result is

$$P(d|x', \mathbb{I}) = \frac{e^{-(d-\bar{d}(x'))^T \left[\Lambda_d^{-1} - \beta^T \beta\right](d-\bar{d}(x'))}}{\sqrt{(2\pi)^{n_d}|\Lambda_d||\Lambda_y||\alpha^T \alpha|}} \ , \tag{8}$$

where $\bar{d}(x')$ is the nominal simulation result $\hat{d}(x', \bar{y})$ and the matrices α and β satisfy the equations

$$\alpha^T \alpha = A^T \Lambda_d^{-1} A + \Lambda_y^{-1}$$
$$\beta^T \alpha = \Lambda_d^{-1} A \ .$$

In this approximation the likelihood is a normal distribution with a modified covariance matrix. The modification is not diagonal, and so the effect of the simulation response is to distort and rotate the likelihood; the rotation is towards an axis determined by the strength of the simulation response to the various nuisance parameters and by the expected variation in those parameters. In the case of current interest, in which experimental data have already been observed, the information associated with the observation d (equation (3)) is modified to become

$$\begin{aligned} I(d|x', \mathbb{I}) = \sum_i \frac{(d_i - \bar{d}(x')_i)^2}{\sigma_i^2} &- \ln P(x|\mathbb{I}) \\ &- (d - \bar{d}(x'))^T \beta^T \beta (d - \bar{d}(x')) \\ &+ \frac{1}{2} \ln \left(|\Lambda_y||\alpha^T \alpha|\right) \end{aligned} \ , \tag{9}$$

which defines a modified χ^2 function that takes into account prior information and the effect of nuisance parameters. These considerations have an important effect on the information associated with a given measurement, and therefore the results of any analysis.

(a) *Nuisance Parameters Ignored*   **(b)** *Nuisance Parameters Included*
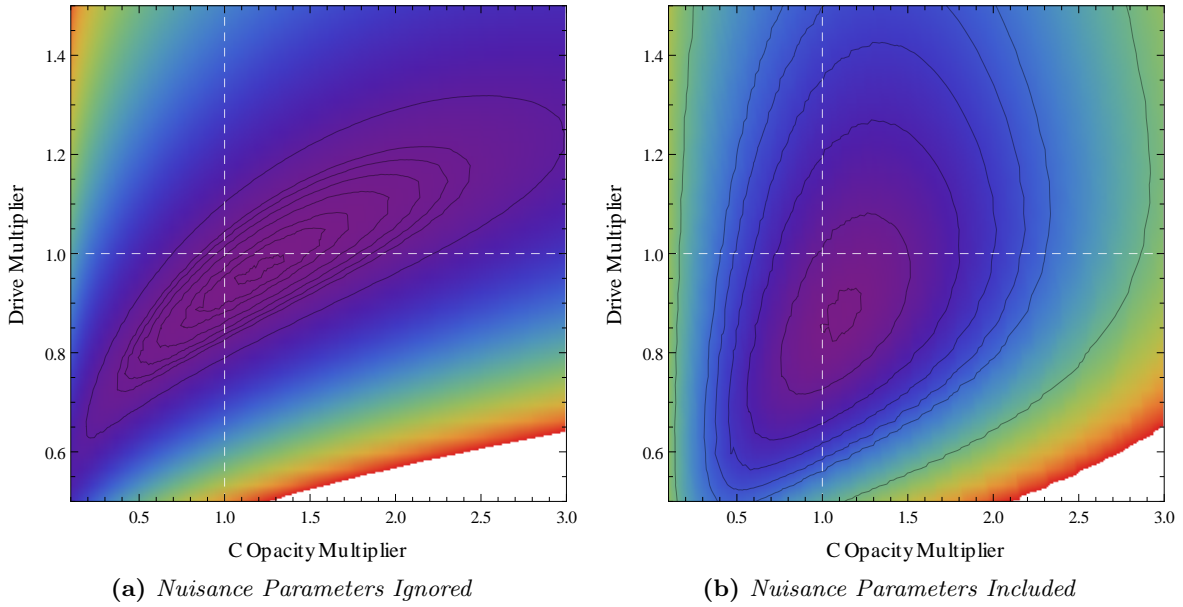
Figure 3: The information contained in the posterior for an analysis of NIF shot N110625. Figure (a) shows the result when variations in the capsule dimensions and atomic composition are neglected, and (b) shows the result when these are included in a linear approximation. Both parameters have a normal prior of standard deviation 0.1. Contour lines show constant fractions of the peak information. Nuisance parameters have the effect of producing a more localised peak in the posterior, however this peak is broader.

4. Application to a NIF Convergent Ablator implosion

The effect of nuisance parameters on N110625 can be included in the above way. We use previous design calculations to populate the response matrix A, treating variations in capsule dimensions and chemical composition as nuisance parameters [15]. These are allowed to vary according to their specified manufacturing tolerances [14]. The modification to the likelihood by these parameters is quite significant; the power of the linear response model is that once the modifications (in the form of the new covariance matrix) are known, subsequent analyses do not need to explicitly include the effect of nuisance parameters at all.

Given knowledge of the prior, from the work in the previous sections and the results of HYDRA simulations, we can infer the drive and carbon opacity multipliers from the described data, including prior information and the effect of nuisance parameters. The results are shown in figure 3; (a) shows the information (equivalently, the logarithm of the posterior) where variations in capsule dimensions and atomic composition are neglected, and (b) show the results when they are included. Both plots use an uncorrelated normal distribution for the prior, with a standard deviation of 0.1 in both directions; as such both represent a more advanced analysis than the χ^2 approach shown in figure 2. The distortion of the likelihood by nuisance parameters has the effect that the peak in the information is more localised, however the overall distribution is broader suggesting a larger error bar on the inferred values.

This can be made more clear by re-casting the data as the posterior probability of the value of the carbon opacity multiplier, by integrating over the drive multiplier. This is not equivalent to the treatment of the drive multiplier as 'nuisance'; the detailed response of the system to changes in drive has been included (not the linear response). This approach will prove more accurate for some rapidly varying parameters, however represents a calculation time penalty. The advantage is in the inspection of correlations between parameters that the multi-dimensional analysis allows. The results of integrating over variations in X ray drive are shown in figure 4, with the blue line showing the inference when nuisance parameters are included, and purple showing the results when they are neglected. There is a significant change in the posterior probability with a corresponding change in the inferred carbon opacity modifier. Table 1 shows the results of this inference for the different Bayesian models we have discussed. These results are (weakly) sensitive to the choice of prior; in cases where the priors are taken as centred on 1, the inferred carbon opacity multiplier is

Inference Model	C opacity Multiplier
χ^2	$1.79^{+0.79}_{-0.46}$
+ prior	$1.23^{+0.23}_{-0.25}$
+ nuisance parameters.	$1.13^{+0.26}_{-0.25}$

Table 1: The results of the inference of the carbon multiplier from measured values of the implosion velocity, radius and fraction of ablator remaining in NIF convergent ablator shot N110625. Error bars are 68% confidence limits. Inference is based on a database of 1024 post-shot HYDRA simulations, using the Bayesian inference models described in the text. The inclusion of prior information and of the variation in capsule dimensions and chemical composition is shown to signficantly reduce the measured carbon opacity multiplier, and the error bar on the result.
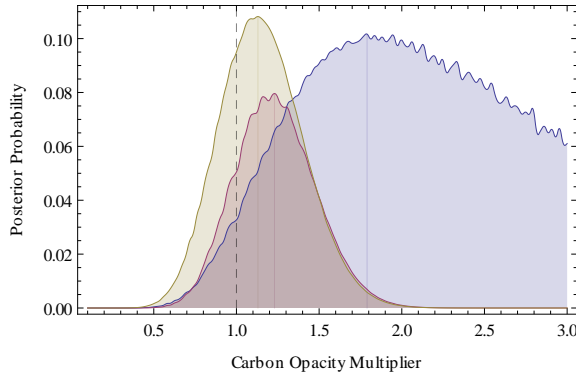

Figure 4: The posterior probability of the carbon opacity multiplier, calculated neglecting nuisance parameters (purple), and including them (yellow). Also shown are the results when a standard χ^2 analysis is used (blue). Variations in target properties cause a shift in the maximum likelihood estimate of the opacity multiplier, and a change in the shape of the distribution of values.

never larger than 1.2.

5. Analysis of a High-Dimensional Parameter Space

The complexity of ICF experiments, and of the hydrodynamic simulations used to analyse and design them, make it unfeasible to perform the simple grid parameter scans that have concerned us until now. The statistical framework does, however, allow more advance Monte-Carlo approaches to be applied, and these have seen great success in the analysis of other large experiments. The application to ICF experiments will still require careful consideration, however, and in the following we will discuss the possible approaches that are available.

5.1. Linear Response

The linear response model described above allows the inclusion of nuisance parameters as described above, however the dependance on parameters that are of experimental interest should be treated more accurately. There is an advantage to keeping as many simulation parameters as possible out of the linear response model; as we have seen in the previous low dimensional analysis, inspection of the posterior as a function of the model parameters can provide valuable insight. It is not possible to include all parameters however, and so the linear response model will be essential in the implementation of the following advanced methods.

5.2. Markov Chain Monte Carlo

The MCMC approach is a statstical method of generating a set of samples of parameter space such that the distribution of samples reflects a given, unknown, probability distribution [10]. The method has been used extensively to probe posterior distributions with very large numbers of parameters; the potential pitfall is that using HYDRA simulations to form the step probability in these models, even using the efficient routines applied here, presents a significant computational challenge.
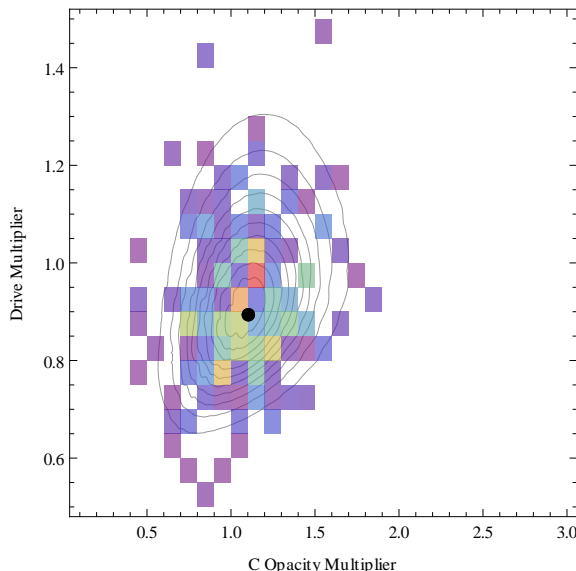
Figure 5: Results of an MCMC simulation of the posterior distribution of carbon opacity and drive multipliers for NIC shot N110625. Contours show the analytically derived posterior, and colored blocks show a histogram of the MCMC results. Nuisance parameters are included using a linear response likelihood, and HYDRA results are interpolated from the simulation grid described in the text.

The feasability of the MCMC approach can be investigated using a surrogate model, that will approximately reproduce the multi-dimensional behaviour of HYDRA simulations. In this way the number of steps required can be found, before resorting to time-consuming HYDRA simulations. An example solution is shown in figure 5, where we plot a histogram of MCMC samples from a run of 2000 steps (initialised at the peak of the prior and with a 500 step 'burn in' period). These simulations use the Metropolis-Hastings algorithm with a normally distributed jump probability to explore the posterior probability formed from the linear response likelihood and prior described in previous sections. The likelihood is approximated by interpolating in the simulation grid described previously, with very good results.

It is well known [3] that even when MCMC simulations poorly describe the multi-dimensional target distribution, low dimensional quantities are often well described. This is seen in our results where the MCMC average and standard deviation carbon opacity multiplier are 1.10 ± 0.23; for an MCMC run with 1000 (post burn in) samples, this value is 1.09 ± 0.22 which still shows good accuracy.

The very small number of simulations required to give a good reproduction of the opacity multiplier is encouraging. It should be noted, however, that even though 2500 HYDRA simulations are easily acheivable using the parrallel control routines used here, MCMC simulations do not lend themselves to easy parrallelisation. The genetic algorithm approach described in the next section are more ameanable to the large scale simulations necessary here.

5.3. Genetic Algorithms

A GA is a computational tool where a set of samples of parameter space are selected and iterated in order to find the extreme value of some merit function, or fitness. New samples are generated from the old using methods inspired by natural selection - those members that have a large fitness are preferentially selected, and bred together. Random jitter is introduced to avoid local maxima in the global fitness, mimiking random mutation in the current population. In data analysis applications it is usual to use the inverse χ^2 function as a fitness, and in this application genetic algorithms are a very robust tool for finding the best fit to data when the parameter space has a high dimension. In previous sections we noted that the usual χ^2 function can be interpreted as the information in the posterior, and showed how this can be modified to include variations in nuisance parameters and prior information. These considerations allow the application of genetic algorithm techniques to our Bayesian analysis.

In figure 6 we plot the evolution of the mean carbon opacity and drive multipliers of a population of 100 parameter space samples as subsequent generations are generated. We use a simple GA, where half the population is carried over between generations and the entire population is allowed to mutate. The breeding population is chosen according to the inverse of the generalised χ^2 discussed previously. The carbon opacity approaches the expected value of 1.10 quickly,
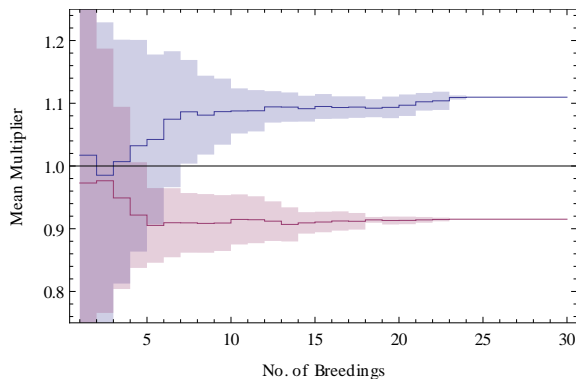
Figure 6: The population mean and standard deviation carbon opacity (blue) and drive (red) multipliers as a function of genetic algorithm breedings. In this simulation a population of 100 simulations were used at each step; this graph therefore represents 2000 HYDRA simulations, which were in this case simulated using an interpolation of a grid of simulations

a value that is quite robust as the number of population members (and therefore HYDRA calculations) is reduced. In the case considered here a full GA simulation would require less than 2000 HYDRA simulations, representing a modest computational problem.

The genetic algorithm technique is very easily parrallelised since all members of the sample population are known at the start of each iteration. It has the disadvantage, however, that the converged result gives no information regarding the form of the posterior distribution and so it is difficult to provide confidence intervals for the inferred parameters, or the information entropy in the posterior. A simple exploration of of the space close to the maximum, in order to calculate the Hessian matrix [3], would provide an efficient solution to the problem.

6. Discussion and Conclusions

We have introduced the Bayesian method and developed the normal linear model to describe nuisance parameters in integrated high energy density experiments. This approach allows a modified Bayesian χ^2 function to be defined that efficiently includes prior information and nuisance parameters. This function can be easily applied in any data analysis or experimental design problem, and is particularly well suited to systems with a large number of nuisance parameters. The compact format provided by the modified covariance matrix means that once a sensitivity study has been performed for a given experimental design, the results can be encorporated in all subsequent data analyses in a simple and consistent manner. The expression of the likelihood in these terms also allows powerful Bayesian and Information-Theoretic results to be used.

These considerations have been applied to a NIF convergent ablator experiment, where it has been shown that they lead to a significant change in the inferred values of the carbon opacity. It has also been shown that the use of computational methods to speed up analysis is very feasible. This capability, along with the Bayesian aspects presented here, will allow the integrate analysis of entire experimental campaigns at the NIF and to the staistical design of future experiments.

7. Acknowledgments

The authors would like to thank Dan Clark for providing design calculations and useful discussion, Robbie Scott and John Castor for help with HYDRA simulations, and Vijay Sonnad for stimulating comments.

References

[1] EL Moses. The national ignition facility and the national ignition campaign. *IEEE Transactions on Plasma Science*, 38(4):684–689, 2010.

[2] SH Glenzer *et Al.* First implosion experiments with cryogenic thermonuclear fuel on the national ignition facility. *Plasma Physics and Controlled Fusion*, 54(4):045013, 2012.

[3] DS Sivia. *Data Analysis: A Bayesian Tutorial*. Clarendon Press, 1996.

[4] KM Hanson. Inference about the plastic behaviour of materials from experimental data. LANL Report LA-UR-04-3529.

[5] D Dodt *et Al.* Reconstruction of an electron energy distribution function using integrated data analysis. *Journal of Physics D: Applied Physics*, 41:205207, 2008.

[6] AM Broomhall *et Al.* A comparison of frequentist and bayesian inference: searching for low-frequency p modes and g modes in sun-as-a-star data. *Monthly Notices of the Royal Astronomical Society*, 406:767–781, 2010.

[7] J Sacks *et Al.* Design and analysis of computer experiments. *Statistical Science*, 4(4):409–435, 1989.

[8] K Chaloner and I Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.

[9] H Dreier *et Al.* Bayesian experimental design of a multichannel interferometer for wendelstein 7-x. *Review of Scientific Instruments*, 79:10E712, 2008.

[10] PJ Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732, 1995.

[11] IE Golovkin *et Al.* Multi-criteria search and optimization: an application to x-ray plasma spectoscopy. *Evolutionary Computation*, 2:1521–1527, 2000.

[12] B Williams *et Al.* Combining experimental data and computer simulations, with an application to flyer plate experiments. *Bayesian Analysis*, 1(4):765–792, 2006.

[13] A Zellner. Optimal information processing and Bayes's theorem. *The American Statistician*, 42(4):278–281, 1988.

[14] SW Haan *et Al.* Point design targets, specifications, and requirements for the 2010 ignition campaign on the national ignition facility. *Physics of Plasmas*, 18:051001, 2011.

[15] D Clark. *Private Communication*

[16] D Hicks. *Private Communication*