

Optimizing the configuration of plasma radiation detectors in the presence of uncertain instrument response and inadequate physics

P.F. Knapp^{1,†}, W.E. Lewis¹, V.R. Joseph², C.A. Jennings¹ and M.E. Glinsky^{1,‡}

¹Sandia National Laboratories, Albuquerque, NM 87185, USA

²Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

(Received 31 May 2022; revised 29 November 2022; accepted 30 November 2022)

We present a general method for optimizing the configuration of an experimental diagnostic to minimize uncertainty and bias in inferred quantities from experimental data. The method relies on Bayesian inference to sample the posterior using a physical model of the experiment and instrument. The mean squared error (MSE) of posterior samples relative to true values obtained from a high fidelity model (HFM) across multiple configurations is used as the optimization metric. The method is demonstrated on a common problem in dense plasma research, the use of radiation detectors to estimate physical properties of the plasma. We optimize a set of filtered photoconducting diamond detectors to minimize the MSE in the inferred X-ray spectrum, from which we can derive quantities like the electron temperature. In the optimization we self-consistently account for uncertainties in the instrument response with appropriate prior probabilities. We also develop a penalty term, acting as a soft constraint on the optimization, to produce results that avoid negative instrumental effects. We show results of the optimization and compare with two other reference instrument configurations to demonstrate the improvement. The MSE with respect to the total inferred X-ray spectrum is reduced by more than an order of magnitude using our optimized configuration compared with the two reference cases. We also extract multiple other quantities from the inference and compare with the HFM, showing an overall improvement in multiple inferred quantities like the electron temperature, the peak in the X-ray spectrum and the total radiated energy.

Key words: fusion plasma, plasma diagnostics

1. Introduction

A common task when studying plasmas in the laboratory is the measurement and interpretation of X-ray radiation. X-ray emission is particularly useful in diagnosing high energy density physics (Matzen *et al.* 2005; Drake 2006) and inertial confinement fusion (Nuckolls *et al.* 1972; Lindl 1995) experiments where the spectral characteristics and

[†] Email address for correspondence: pfnapp@sandia.gov

[‡] Present address: qiTech Consulting, Santa Fe, NM, USA.

total output of the X-ray emission can help to constrain important physical quantities like pressure, temperature and the areal density of the confining shell (Ma *et al.* 2013; Knapp *et al.* 2019). Additionally, when studying the effects of radiation on materials and electronics, these detectors are crucial for understanding the total output and spectral content of the radiation source (Coverdale *et al.* 2010; Ampleford *et al.* 2014). In magnetic confinement fusion devices, such as DIII-D (Luxon 2002), JET (Rebut, Bickerton & Keen 1985; Keilhacker & Team 1999) and ITER (Iter 1999; Aymar, Barabaschi & Shimomura 2002), arrays of soft X-ray detectors are used to measure spatially resolved temperatures (Alper *et al.* 1997; Delgado-Aparicio *et al.* 2007) which help to constrain the configuration and performance of the plasma.

Unfortunately, the detectors used to measure the radiation of interest are difficult to calibrate resulting in large uncertainties and must be fielded in harsh environments that contribute artefacts to the data. Furthermore, the detectors integrate over photon energy and space, meaning that very little of the spectral information is preserved in the raw measurement, making the extraction of important physical quantities an ill-posed inverse problem. In order to extract useful physical information from a set of radiation detectors we must impose a model on the analysis that relates observed diagnostic signatures to the parameters of interest. The forward physics model used to do this is quite often an oversimplification of the object under study, necessitated by computational convenience and/or interpretability. These simplifications can introduce bias in the inferences made from measurements and must be understood. Finally, the precision with which we can infer a specific quantity of interest (QOI) will depend not only on the quality of the detector calibration, but also on our choices as experimentalists regarding how to configure the instrument.

Bayesian inference is a popular choice of methodology to solve the inverse problem of inferring QOIs from measurements where prior information is used to regularize the solution and probabilistic models are used to sample the posterior distribution providing the experimentalist with most-likely parameter values as well as credibility intervals and correlations (Wikle & Berliner 2007; Von Toussaint 2011). Uncertainties in instrument responses can be incorporated through prior distributions on their values which express our degree of belief in a certain value before the observations are made.

Here, we demonstrate an experimental design methodology using synthetic experiments that takes advantage of Bayesian inference for uncertainty quantification, and optimizes the configuration of a set of radiation detectors to maximize the confidence in our inferences within the available resources. We show that, through proper choice of optimization metric, this procedure simultaneously minimizes the uncertainty and the bias in the inference, allowing us to optimally configure our instruments to provide unbiased estimates to the extent possible in the presence of an inadequate forward physics model and uncertain instrument response characteristics.

The remainder of this paper is organized as follows. In § 2 we define the inference and optimization problem of interest, identifying the general properties of each step and the form of the optimization metric. In § 3 we apply this framework to the problem of interest, which is intended as a pedagogical example with realistic properties, identifying the models used and parameters needed. We further develop the metric, showing different options for assessing error and adding a penalty term to deal with real diagnostic limitations that our model cannot naturally take into account. In § 4 we present the results and discuss the performance of our approach. Finally, in § 5 we discuss the application of this method to less idealized problems as well as extensions to this formalism.

2. Problem description

2.1. Inference

In our problem of interest we cannot measure the parameters θ describing the physics of interest directly. Instead we make measurements using instruments that depend implicitly on the values of θ through the physics model. Often we have multiple instruments whose observations have sensitivity to multiple different parameters. Thus using all instruments simultaneously is the most effective way to determine θ that best describes the ensemble of measurements.

We start with a parameterized model of our physical system of interest

$$y = f(\mathbf{d}; \theta), \quad (2.1)$$

where $\mathbf{d}$ are the independent coordinates of interest (e.g. time, spatial coordinates, frequency, etc.), and θ are the parameters. The output of the function f is expressed over the independent coordinates and will be used by diagnostic models to produce synthetic observations. We obtain several diagnostic measures O_i , $i = 1, \dots, M$ which are functions of the output of the physical model and configuration details of the instrument

$$O_i = g_i(y; Z_i) + \delta_i, \quad (2.2)$$

where $g_i(\cdot)$ values are known functions, Z_i values are the known configurations of the diagnostic instruments and δ_i values are the unknown measurement errors.

The Z_i values represent all the information needed to field and interpret the instrumental data. These could be specific choices made when configuring the instrument (e.g. the source-to-detector distance, detector type, probe laser power, filters, attenuation on an oscilloscope, etc.), as well as quantities that define the response characteristics of the instrument (e.g. sensitivity of the detector, thermocouple temperature coefficient, collection solid angle, etc.). Most often the quantities representing specific choices are known with a high degree of certainty and represent a finite set of available configurations. These quantities can be used to control how sensitive a given instrument is to a given QOI. Different instrument configurations may be ideally suited for different experimental configurations and different QOIs, and so choices must be made on a per-experiment basis. Additionally, these choices can also effect the expected signal levels, so background and noise contributions, as well as instrument dynamic range, must be well understood in order to configure an instrument such that a reliable signal is obtained. Quantities in this category are generally not treated as random variables in our formalism, but rather specific choices that must be made.

The quantities representing response characteristics are usually quantities that require calibration and are known with some uncertainty. There can also be unknown bias in the calibration data, or drifts with time, that must be accounted for. As such, these quantities are treated as random variables in the formalism in order to capture the uncertain nature of their values.

Since some instrument parameters are to be treated as random variables, and others not, we will break the Z_i values into two groups: $Z_i = [z_i, \xi_i]$, where the z_i values are deterministic variables, and ξ_i values are random variables. Specifically, the ξ_i values will be treated as normally distributed variables where the mean and standard deviation are known, $\xi_i \sim \mathcal{N}(\mu_i, \sigma_i)$. These uncertain quantities are marginalized out when the posterior is sampled.

Using Bayesian inference we wish to estimate the posterior distribution of the parameters θ , $p(\theta | O)$, where $O = (O_1, \dots, O_M)$. In addition to estimating the model




FIGURE 1. Network describing our experimental inference problem. Model parameters θ with appropriate prior distributions are fed into the forward physics model $f(\mathbf{d}; \theta)$ producing deterministic output quantity y . Model output is fed into the diagnostic models whose behaviours are controlled by the configuration choices z_i and stochastic calibration values ξ_i . The output of the diagnostic models are compared with experimental observations O_i . Additional unobserved quantities of interest are computed as $Y = h(y)$.

parameters θ , we may also wish to estimate some unobserved output quantity $Y = h(y)$. This inference is shown as a network graph in [figure 1](#). Once we establish appropriate prior distributions on the θ and ξ_i , we can use standard Bayesian computational algorithms to sample from the posterior.

2.2. Optimization methodology

Having established a means to infer physical parameters θ from a set of observations $\{O_i\}$, the question then becomes how best to configure our instruments in the face of uncertain response information and simplified physics model to minimize the resulting uncertainty and bias on inferred quantities. This task, as with most instrument design and sensitivity analysis tasks, will be done with synthetic data taken from a high fidelity model (HFM). This approach means we know the true θ values *a priori*, but saying this belies the true complexity of the situation. The HFM will produce a rich variety of data that we cannot possibly hope to infer with our reduced model. This requires that we develop a physics-motivated mapping from the HFM to a reduced representation that we can compare with our model.

Once we have chosen this mapping, we must optimize the configuration of our instrument with a metric that embeds this mapping. This, in principle, is not so challenging an optimization problem, although the parameter space describing an instrument configuration can be quite large. However, when considering real instruments, there is almost always a finite set of configurations that are practical to achieve. While this does restrict the space, it means that the optimization procedure must be able to handle mixed continuous and discrete parameters. Additionally, for real problems the physics forward model and diagnostic models can be quite expensive to run, so choosing an optimization procedure that minimizes the number of samples drawn is advantageous.

For these reasons we apply Bayesian optimization (BO) to the problem (Jones, Schonlau & Welch 1998; Shahriari *et al.* 2016; Frazier 2018). BO allows for optimization of expensive black box functions without access to gradients, it is able to handle mixed parameters (e.g. continuous, discrete, categorical), it readily applies constraints, and it

is efficient at finding suitable optima. BO works by approximating the objective surface using a Gaussian process (GP) which provides an estimate of the mean and variance of the function at all points in the parameter space (Williams & Rasmussen 2006; Schulz, Speekenbrink & Krause 2018). An initial random set of function evaluations is used to create a first guess at the objective surface. The key novelty of BO is that both the mean and the variance are used to choose the next evaluation point. Instead of just finding the maximum value predicted by the GP, the GP is fed into another function, called the acquisition function whose maximum is a compromise between maximizing the mean and maximizing the variance of the GP. This allows the optimization algorithm to balance exploitation and exploration, providing a tendency to search the space and not settle for the first local optimum found. Other optimization algorithms could easily be used instead of BO. This may be necessary as the dimensionality of the problem increases since GPs could become computationally prohibitive in this regime.

Many metrics exist in the literature to help one minimize uncertainty in inferred quantities, such as the Fisher information, etc. (Silvey 1980). Due to computational costs most of the available metrics rely on maximum likelihood or maximum *a posteriori* estimates to form the metric. These methods can introduce bias because they rely on finding an optimum which, particularly in high-dimensional problems, can be far from the highest density portion of the distribution. The information matrix is then estimated using the Hessian at this optimum, which may not represent the curvature of the high density region (Murphy 2022). Since we will be sampling the posterior in our inference of parameters θ we would prefer to retain the information contained in the posterior when computing our optimization metric. Some metrics have been developed to estimate the information gain in the presence of new data utilizing the full posterior, e.g. the Kullback–Leibler divergence (Bishop & Nasrabadi 2006), however, estimating this quantity using samples from a posterior can be a numerically challenging and unstable problem, often necessitating the use of analytic approximations of the posterior, e.g. variational inference (Bishop & Nasrabadi 2006). Here, we propose the use of the mean squared error (MSE) between the distribution of posterior values of q and the true value Q from the HFM given by

$$\text{MSE}^j = \frac{1}{N} \sum_{k=1}^N (q_k^j - Q^j)^2, \quad (2.3)$$

$$\begin{aligned} &= \frac{1}{N} \sum_{k=1}^N (q_k^j + \langle q \rangle^j - \langle q \rangle^j - Q^j)^2, \\ &= (\langle q \rangle^j - Q^j)^2 + \frac{1}{N} \sum_{k=1}^N (q_k^j - \langle q \rangle^j)^2, \end{aligned} \quad (2.4)$$

where N is the number of posterior samples. Here, $q(Q)$ is any quantity of interest that can be derived from the LFM (HFM). This allows us to not only consider the quality of the fit to the model parameters θ , but any latent quantity that is important but not directly measurable. We have added an index j to the quantities in (2.3) to note that in general we will be measuring the quality of the fit against more than one realization of the HFM to avoid overfitting to a single instance representing a single set of plasma conditions. We refer to the set of instances we optimize against as the training data, and each individual instance j as a training sample. In (2.4) we show that the MSE^j is equivalent to the sum of the square of the bias between the mean of q^j conditioned on our observations $\langle q \rangle^j = \mathbb{E}[p^j(q | O_i)]$ and the truth value Q^j computed from the HFM, and the variance

in the posterior $p^j(q | O_i)$ (Bishop & Nasrabadi 2006). Therefore, minimizing the MSE conditioned on the observations O_i over diagnostic configuration z_i ensures that we pick a configuration that will allow an inference maximally consistent with the HFM quantity of choice Q^j across the training samples.

One subtlety regarding this metric is that it is non-negative. Gaussian processes have known deficiencies in approximating these kinds of functions. Additionally, the vanishing value of the MSE as it approaches zero serves to de-emphasize improvements made once it has reached a sufficiently small value. As such, we will use $\log(\text{MSE})$ in our optimization metric, making the function attain both positive and negative values and transforming small values to arbitrarily negative ones, improving the performance of both the GP and the optimization. If using a different optimization technique this may not be necessary. Additional information can be added to the optimization metric to enforce constraints or impose domain specific knowledge. The metric will generally take the form

$$\mathcal{M} = \log \left(\frac{1}{K} \sum_{j=1}^K \text{MSE}^j + L^j \right), \quad (2.5)$$

where L^j is a loss, or penalty, term used to add any additional information that is not explicitly accounted for in the posterior and an average is taken over the training samples before the logarithm is taken.

Our optimization algorithm relies on Bayesian inference to compute the metric. Therefore, at each iteration, we must compute the synthetic data from the HFM using the updated diagnostic parameters so that we can obtain our samples q_i from the posterior and compute the MSE. To initialize our optimizer we generate n space-filling samples of z_i in our parameter space. Our algorithm is summarized as follows:

Algorithm 1 Instrument Optimization

- 1: Generate n random samples of z_i
 - 2: **for all** z_i **do**
 - 3: **for all** training data $j = 1 \dots K$ **do**
 - 4: Generate $O_i^j(z_i, \xi_i)$
 - 5: Sample posterior $p^j(\theta | O_1, \dots, O_M)$ to obtain q_k^j values
 - 6: Compute MSE^j and L^j
 - 7: Compute $\mathcal{M}$
 - 8: Fit GP to initial n points
 - 9: compute expected improvement
 - 10: select new design point z_i to maximize expected improvement
 - 11: **repeat**
 - 12: **for all** training data $j = 1 \dots K$ **do**
 - 13: Generate $O_i^j(z_i, \xi_i)$
 - 14: Sample posterior $p^j(\theta | O_1, \dots, O_M)$ to obtain q_k^j values
 - 15: Compute MSE^j and L^j
 - 16: compute $\mathcal{M}$
 - 17: fit GP to all $\mathcal{M}$ values
 - 18: compute expected improvement
 - 19: select new design point z_i to maximize expected improvement
 - 20: **until** max iteration
-

The iteration stops when a specified stopping criterion is met, typically a maximum number of iterations. To perform our optimization we use the GPyOpt package (GPyOpt authors 2016), which we found to be suitable for our application. In GPyOpt, discrete variables are handled by marginally optimizing over feasible values, which can be extremely slow if many discrete variables are used, but is tractable in our case. Other packages exist that may be able to better take advantage of modern computing architectures, e.g. GPUs, for improved computational efficiency and parallelism (Knudde *et al.* 2017; Balandat *et al.* 2020).

Before moving on, we note that steps 8–9 and 17–18 in our design optimization algorithm constitute the use of BO. In principle, BO could be replaced with any other appropriate optimization algorithm given the problem structure. For example, genetic algorithms (Mitchell 1998) could be utilized for problems with mixed discrete and continuous variables as demonstrated here. For continuous variable problems, gradient based methods may be appropriate. We refer the reader to the textbook (Kochenderfer & Wheeler 2019) and references therein for a review of available methods.

3. Radiation detector optimization

As an exemplar problem, we will explore the use of radiation power detectors to measure certain properties of hot, dense plasmas. In experiments that generate hot dense plasmas, such as in magnetic confinement fusion devices, inertial confinement fusion plasmas and other similar objects, intense X-ray radiation is generated from the object under study. The X-ray emission from these objects carries with it information about the state of the plasma, including its temperature, pressure and volume. A simple diagnostic used to measure this emission is an array of radiation power detectors. The exact kind of detector varies depending on the characteristics of the experiment. We will consider photoconducting diamond (PCD) detectors, which are pieces of diamond over which a bias voltage is applied, typically ~ 100 V. When X-ray photons are absorbed in the diamond a photocurrent is produced, which is proportional to the X-ray power incident upon the detector. The response of the PCD is therefore determined by the frequency dependent absorption of X-rays in the detector and the sensitivity of the element.

Often, numerous PCDs are fielded on a given experiment with a variety of X-ray filters in front of each element. The filters attenuate different portions of the X-ray spectrum, producing weighted integrals over photon energy. These weighted integrals are what allow us to obtain information about properties of the emitting plasma. By way of providing a concrete example to study we will consider the emission from a cylindrical deuterium plasma surrounded by a beryllium liner. This configuration is similar to the plasmas studied at the stagnation phase of MagLIF experiments (Slutz *et al.* 2010; Gomez *et al.* 2015, 2020; Yager-Elorriaga *et al.* 2021) fielded on the Z machine at Sandia National Laboratories (Savage *et al.* 2011).

In order to test our inferences and optimize the diagnostic configuration we must first generate synthetic data from our HFM. We are using an ensemble of one-dimensional simulations implementing the GORGON magneto hydrodynamics (MHD) approach (Chittenden *et al.* 2004; Ciardi *et al.* 2007; Jennings *et al.* 2010) to provide these data. From each calculation we post-process the data to produce the spatially integrated spectrally resolved emitted power, given by

$$P_{\epsilon}(t) = \int_{V(t)} \exp(-\tau_{\epsilon}^{\ell}) \mathcal{E}(\mathbf{r}, t) dV, \quad (3.1)$$

where $\varepsilon(\mathbf{r}, t)$ is the plasma emissivity as a function of space and time and $\exp(-\tau_\epsilon^\ell)$ accounts for attenuation from the liner. The emissivity is computed using a bremsstrahlung emission model which takes the density and temperature of the plasma at each point in space and time as inputs (Knapp *et al.* 2019, 2022). The radiated power is the quantity that the PCDs are directly sensitive to, so we can use this to produce synthetic voltage signals from a suite of PCDs. This signal depends on the input impedance of the oscilloscope I (assumed to be 50Ω), the sensitivity of the element S , the solid angle subtended by the element $\Delta\Omega$, the spectral absorptivity of the detector and the transmission of the filter applied to the detector. The solid angle is $\Delta\Omega = 4\pi d^2/A$, where d is the distance from source to detector and A is the active area of the detector. In general, the transmission of X-rays through a material can be written as

$$T_\epsilon = \exp(-\rho\ell\kappa_\epsilon), \quad (3.2)$$

where ρ is the density of the material, ℓ is the path length of the X-rays through the material and κ_ϵ is the photon energy-dependent opacity. It follows that the absorption is just $A_\epsilon = 1 - T_\epsilon$. Therefore, we can write the signal observed on the oscilloscope as

$$o(t) = \frac{SI}{\Delta\Omega} \int_0^\infty d\epsilon T_{\epsilon, \text{filter}} A_{\epsilon, \text{PCD}} P_\epsilon(t). \quad (3.3)$$

Finally, we add Gaussian noise with a mean of 0 and a standard deviation of 50 mV to the synthetic observations. When generating synthetic data we add a random bias to the detector sensitivity, sampled from a normal distribution with zero mean and standard deviation equal to the uncertainty in the detector response, to account for the fact that the true sensitivity is poorly understood. Given this model of the synthetic observations we have multiple free parameters to choose. In order to restrict this set we fix some of these values to those typically used on relevant experiments on Z (Jones *et al.* 2014). The source to detector distance is 170 cm, the detector area is 0.01 cm^2 and the detector thickness is 0.05 cm. Fixing these parameters allows us to fix the detector absorption and solid angle. These values are known with a small uncertainty in practice, so we use a prior with fixed mean and variance to account for any small uncertainty we have in their true values. This leaves three remaining quantities that need to be chosen in order to fully specify the response of each detector, namely the filter material m , filter thickness δ and detector sensitivity S .

In order to interpret these signals we must develop a simple model of the X-ray emission that is efficient enough to be used in Bayesian inference, yet retains enough of the necessary physics to be meaningful. We refer to this as our low fidelity model (LFM). For our LFM we assume a spatially and temporally uniform pure deuterium plasma surrounded by a beryllium liner. The X-ray emission as a function of photon energy is written as

$$Y_\epsilon = 2\Delta t A_{ff} V_{HS} P_{HS}^2 \exp(-\tau_\epsilon^\ell) \frac{g_{ff} Z}{(1+Z)^2} j, \quad (3.4)$$

where

$$j = \left(Z^2 + \frac{A_{fb} \exp(RyZ^4/T_e)}{A_{ff} T_e} \right) \frac{\exp(-\epsilon/T_e)}{T_e^{5/2}}, \quad (3.5)$$

is an analytic approximation for the emissivity of hydrogen as a function of photon energy ϵ , with the Rydberg constant $Ry = 13.6 \text{ eV}$, the free-free A_{ff} and free-bound A_{fb} emission coefficients, and $g_{ff} = 2(0.87\sqrt{3}/\pi)\sqrt{T_e/\epsilon}$ is the free-free Gaunt factor (Epstein *et al.*

2015). The deuterium fuel has volume V_{HS} and is at pressure P_{HS} and temperature T_e and exists for duration Δt . Since we are restricting ourselves to a pure deuterium plasma $Z \equiv 1$. Finally, the term $\exp(-\tau_e^\ell)$ captures the attenuation of X-rays emitted from the fuel as they pass through the beryllium liner before reaching the detector, where the liner optical depth is defined as $\tau_e^\ell = \rho R_\ell \kappa_{\epsilon, Be}$. The attenuation is governed by the beryllium opacity $\kappa_{\epsilon, Be}$, and the areal density of the liner ρR_ℓ .

The only parameters in this model that affect the shape of the spectrum are the temperature T_e and liner areal density ρR_ℓ . Therefore, for purposes of illustration we will concern ourselves with estimating the plasma temperature T_e , liner areal density ρR_ℓ , a constant scale factor, which is a nuisance parameter, and the total integrated output energy. We can simplify the expression above, and since the emission is uniform in time, fold the emission duration Δt into the scale factor to obtain the spectrally resolved radiated energy

$$Y_\epsilon = C \exp(-\tau_e^\ell) \frac{1}{4} g_{ff} j. \quad (3.6)$$

Additionally, we obtain the total radiated energy as $Y = \int_0^\infty y_\epsilon d\epsilon$.

It is important to note that there is an obvious mismatch between the physics contained in our synthetic data and our LFM. The synthetic data are time dependent, but our simple model is constant in time and provides us with a radiated energy, not power. Therefore, if we plug (3.6) into (3.3) we will obtain a signal in units of $V \cdot s$, not V . We cannot easily add time evolution to our model, so the simplest way to overcome this is to integrate (3.3) in time allowing us to compare our model directly with the synthetic observation.

We have now constructed a model that allows us to directly compare synthetic data from our HFM with data generated from our LFM for a given set of parameters. Recasting this in the notation developed in § 2.1 we have $\theta = \{T_e, \rho R_\ell, C\}$, $z_i = \{m_i, \delta_i, S_i\}$, and $\xi_i = \{A_i, d_i\}$:

$$f(\epsilon; \theta) = C \exp(-\rho R_\ell \kappa_\epsilon) \frac{1}{4} g_{ff} j \quad (3.7)$$

$$g_i(y; Z_i) = \frac{S_i I}{\Delta \Omega} \int_0^\infty d\epsilon T_{\epsilon, i} A_{\epsilon, \text{PCD}} f(\epsilon; \theta) \quad (3.8)$$

$$h(y) = \int d\epsilon f(\epsilon; \theta). \quad (3.9)$$

Now we may specify our log-likelihood function, for which we choose a normal distribution, giving

$$\log \mathcal{L} \propto -\frac{1}{2} \sum_i \frac{(g_i - O_i)^2}{\sigma_i^2} - \sum_i \log(\sqrt{2\pi} \sigma_i). \quad (3.10)$$

Combining this with appropriate priors for all parameters we obtain a statistical model that can be sampled with standard algorithms. For the temperature we set $p(T_e) = \mathcal{N}(4, 3 \text{ keV})$, with bounds placed at 0.5 and 10 keV. For ρR_ℓ we set $p(\rho R_\ell) = \mathcal{N}(1, 1 \text{ g cm}^{-2})$, with bounds at 0.1 and 3 g cm^{-2} . The bounds for temperature and ρR_ℓ are set to fully encompass reasonable stagnation states, but limit unrealistic conditions. In our model, the scale is extracted from the HFM by measuring the volume, pressure and duration of stagnation such that C is expected to be $O(1)$. Accordingly, we use a normal distribution for $\log_{10}(C)$, $p(\log_{10}(C)) = \mathcal{N}(\log_{10}(1), \log_{10}(3))$ with bounds at $\log_{10}(0.3)$ and $\log_{10}(30)$ to allow for significant departure from this expectation. Uncertainties on the filter thickness, detector sensitivity, detector distance and detector area are accounted for

using normal priors with standard deviations of 5 %, 5 %, 0.5 % and 0.5 % respectively. We used pymc3 (Patil, Huard & Fonnesbeck 2010) to perform the sampling. We found that acquiring 8000 MCMC samples per inference provided a converged chain with negligible variance in computed expectation values.

In order to find the set of instrument parameters z_i that minimizes the bias between the posterior mean and the ground truth, and the variance in the posterior, we have to decide relative to what. This is a subtle choice, and the so-called ‘truth’ values to which we are comparing can affect the configuration we end up deciding is ‘best’. It is tempting to choose the temperature and liner areal density from the HFM as the reference values. However, as noted earlier, our model takes as its input a single temperature and ρR_ℓ while the HFM model produces values for these quantities that vary in space and time. How does one map from these high-dimensional spaces down to the compact representation we have chosen? The answer is that it depends on what we wish to know and the choice of mapping is not obvious. This ambiguity is a direct result of the missing physics in our LFM. To circumvent this ambiguity we propose to use the time integrated spectrum $Q = Y_\epsilon$ as the metric for comparison. This quantity is rigorously defined from both the HFM and LFM with no ambiguous interpretation. Additionally, it is the fundamental quantity from the experiments that is being observed, whereas quantities like temperature and ρR_ℓ are derived from the spectrum. Once we have a configuration that produces a sufficiently accurate fit to the spectrum we can relate the posterior distribution of model parameters to physical quantities of interest secure in the knowledge that they faithfully represent the spectrum and not some *ad hoc* mapping.

With this choice, the expression for the MSE for a single training sample can be written as follows:

$$\text{MSE}^j = \frac{1}{N} \sum_{k=1}^N \int_0^\infty d\epsilon \left(\frac{y_{\epsilon,k}^j - Y_\epsilon^j}{Y_\epsilon^j} \right)^2, \quad (3.11)$$

where $y_{\epsilon,k}^j$ is the k th sample from the posterior of the emitted spectrum and Y_ϵ^j is the true spectrum from the HFM from the j th training sample. The integral is over photon energy and the summation is over samples of the posterior. We scale the MSE by the true spectrum to allow the emission at all photon energies to be weighted more evenly such that the peak of the spectrum does not dominate the MSE.

There is one final factor we must consider when performing our optimization. In practice, signals are limited in dynamic range by a variety of factors. For our purposes, PCD signals are limited on the low end by noise and on the high end by the bias voltage. A PCD cannot produce a signal larger than the applied bias voltage, but in reality there is a nonlinear compression of the signal that begins at much lower signal amplitudes. The noise limitation will manifest as large variance in the posterior when signals are sufficiently low in amplitude causing large uncertainty in their values. In a sense, the negative aspects of low signal strength will be accounted for naturally in the MSE. However, as a direct result of the form of the LFM we have chosen the bias voltage limitation will not. Our model explicitly assumes a stationary, uniform plasma and produces a total radiated energy, meaning the time-dependent compression of the signal cannot be accounted for in the inference. Adding time dependence to our LFM would require a substantial increase in the complexity of the model and represents a significant challenge. In order to account for this effect we must introduce a loss term that penalizes PCD configurations that produce excessively large signals. This loss term is inherently domain specific and will often come at the expense of hyper-parameters that are not defined *a priori* by any physical or statistical arguments. While this introduces subjectivity to the problem, it also introduces

a means for the experimenter to exercise control over which concerns are emphasized in the final solution.

As stated, the voltage recorded on the oscilloscope is proportional to the incident X-ray power with a nonlinear correction (Spielman, Hsing & Hanson 1988; Jones *et al.* 2014), given as

$$V_{\text{osc}} = \frac{o(t)}{1 + \frac{o(t)}{V_{\text{bias}}}}, \quad (3.12)$$

where $o(t)$ is computed from (3.3) and V_{bias} is the applied bias voltage. We can see that when $o(t) = V_{\text{bias}}/2$ the recorded signal is $\sim 2/3$ the actual signal. Our ability to recover the true signal amplitude from the measured is highly uncertain. Furthermore, as $o(t) \gg V_{\text{bias}}$, the recorded signal will saturate at V_{bias} , making it impossible to recover accurate estimates of the true voltage. For this reason we want to penalize large voltages in our optimization, and we want to penalize them more strongly as they get larger. We therefore choose an exponential form for the penalty term

$$L^j = \exp\left(\frac{\max(V_{\text{peak},i})^j}{\alpha V_{\text{bias}}}\right) - 1. \quad (3.13)$$

In this expression $V_{\text{peak},i}$ is the peak voltage on detector i from (3.3). We take only the largest peak amongst the detectors fielded for each training sample to compute the penalty. This form has two desirable features: (i) it strongly penalizes signals that are larger than αV_{bias} with increasing penalty as V_{peak} increases, and (ii) subtracting 1 gives $L \rightarrow 0$ for values of $V_{\text{peak}} \ll \alpha V_{\text{bias}}$, ensuring that the penalty has no effect on the optimization if signals are kept small. With the form of the penalty chosen, we now have the full expression for the metric we wish to optimize

$$\mathcal{M} = \log\left(\frac{1}{K} \sum_{j=1}^K (\text{MSE}^j + \lambda L^j)\right). \quad (3.14)$$

Where the summation over training samples produces an ensemble average of the MSE and the penalty. Unfortunately, we are left with two hyperparameters, α and λ , that must be chosen empirically. Intuitively, λ controls how strong of an effect the penalty has on the optimization as a whole and α determines the voltage threshold where strong penalization begins to take off. A small hyper parameter scan was conducted to find reasonable values (see Appendix B). Based on this scan we set $\lambda = 0.15$ and $\alpha = 0.25$ for the results that follow.

4. Results

To perform the optimization we draw from a database of one-dimensional simulations of MagLIF implosions, which makes up our HFM. The ensemble comprises a variety of input parameters (e.g. laser energy coupled, liner dimensions, gas density and initial magnetic field strength), producing a range of stagnation conditions and X-ray outputs. The distribution of emissivity-weighted temperature, liner ρR_ℓ and total X-ray output from this ensemble are shown in figure 2. From this ensemble we draw four realizations that span the range of the distribution, shown in magenta, which will be used as training data to perform the optimization. These points are selected using support points (Mak & Joseph 2018; Joseph & Vakayil 2022) to accurately represent the full distribution. As




FIGURE 2. Pairwise distributions of emissivity-weighted temperature and liner areal density as well as total X-ray output from the ensemble of one-dimensional simulations. (a) Distribution of liner areal density with temperature. (b) Distribution of $\log_{10}$ of the X-ray output with temperature. (c) Distribution of $\log_{10}$ of the X-ray output with liner areal density. Grey points show the entire dataset, blue points show those used for validation and magenta show those used for training.




FIGURE 3. (a) Best value of the optimization metric as a function of iteration for runs initialized with five different seeds. The run in green arrived at the best solution. (b) The spectral response of each of the five detectors including the filter for the best solution out of the five runs.

stated, multiple instances of the HFM are chosen to avoid a solution that is too finely tuned to the specific conditions achieved in a single realization, which may not generalize well. Finally, using support points we select another set of 16 instances from the ensemble (shown in blue), which we use as validation data to test the performance of the optimized configuration against two reference configurations. The reason we use such a small number of points to train the diagnostic configuration is due to computational cost as the posterior must be computed for each instance in the training set for each round of the optimization.

Figure 3(a) shows the results of seven separate optimization runs initialized with different random seeds. This plot shows the best value of the optimization metric as a function of iteration, clearly demonstrating the improvement in the optimization metric as a function of iteration. Each run is initialized with 10 Latin hypercube (Mckay, Conover & Beckman 1979) samples from the parameter space. The range of allowed filter thicknesses is [5, 500 μm] and the list of allowed filter materials is given in Appendix A. We can clearly see that each run arrives at a different best configuration, demonstrating the non-convex nature of this problem. For further analysis we use the results of the run shown

Detector	Material	Thickness (μm)	Sensitivity (A/W)
1	Molybdenum	4.7	1×10^{-4}
2	Palladium	18.1	25×10^{-4}
3	Vanadium	26	25×10^{-4}
4	Vanadium	4.2	1×10^{-4}
5	Kapton	45	1×10^{-4}

TABLE 1. Final configuration of filter material, filter thickness and detector sensitivity for each of the 5 PCD elements determined through the optimization.

in pink which arrived at the best overall solution. The spectral response of each detector which includes the transmission through the filter and the absorption in the detector is shown in figure 3(b). The configuration is detailed in table 1, which lists the filter material, thickness, and detector sensitivity for each detector.

In order to assess the quality of this configuration we turn to our set of validation cases from the ensemble (blue circles in figure 2), as well as two different filter sets we will use as reference cases for comparison. The first reference case, referred to as the ‘MagLIF’ configuration, is one that is used typically on MagLIF experiments conducted on the Z machine. This configuration uses only three PCDs with different thicknesses of Kapton (127, 508 and 762 μm) as the filters. The response characteristics are shown in figure 4(a). We note that this configuration was never intended to be used to infer stagnation quantities, only to ensure that a stagnation signal was measured in the face of large uncertainty in the expected signal on the first MagLIF experiments. It is still used because it achieves this goal reliably. For our purposes, we choose the detector sensitivity to balance the three different signal strengths. The second reference configuration, referred to as the ‘Expert’, was arrived at by choosing a set of 5 filters that overlap in ways that provides sensitivity to different portions of the spectrum. This configuration is shown in figure 4(b) where it can be seen that the ensemble of filter edges and bandpass regions provide weighting to different regions of the spectrum. We, again, manually select the detector sensitivities to balance the signal strengths as best as possible. We emphasize that this is a reasonable, if not optimal, filter configuration that was chosen in a manner that might be expected in practice in the absence of a clearly defined optimization metric. Using these configurations, we create the synthetic data from each case in our validation set and run the inference, producing posterior distributions of the spectrum and each of the model parameters.

4.1. Comparison with spectrum

To qualitatively assess the differences in performance between the three configurations we look at the spectra reconstructed with each configuration. Figure 5 shows the posterior spectrum produced for three randomly chosen cases from the validation set for each of the different filter configurations. For each case, the true spectrum is shown as the dashed black line, the median of the posterior is shown as the solid line and the shaded bands show the 95 % and 68 % credible intervals. The top row shows the results for the ‘MagLIF’ configuration, the middle row shows the results for the ‘Expert’ configuration and the bottom row shows the results for the ‘Optimum’ configuration. Looking at the top row, we see that the ‘MagLIF’ configuration performs very poorly on each of the cases chosen. The credible intervals are very large and, although they encompass the true spectrum




FIGURE 4. (a) Spectral response for the ‘MagLIF’ reference configuration. (b) Spectral response for the ‘best guess’ reference case.

(dashed black line), the median spectrum (solid blue) agrees very poorly with the truth. This configuration is clearly not optimal. Examining the middle row, corresponding to the ‘Expert’ configuration, we see a marked improvement. The credible intervals are reduced and the median is closer to the truth in all three cases, although there are still significant discrepancies, particularly at low photon energies. Finally, examining the performance of our ‘Optimum’ (bottom row) we see excellent agreement with the true spectrum in all three cases. The credible intervals are further reduced and the median is very close to the truth. Generally, the shape of the spectrum is very well reproduced.

To better understand systematic differences we can examine the performance of each configuration across the entirety of the validation dataset. We define the difference between the posterior and the true spectrum, normalized to the true spectrum, as $\Delta_\epsilon = (y_\epsilon - Y_\epsilon)/Y_\epsilon$. With this metric, a perfect fit will produce 0 at all photon energies, positive values where the posterior is larger than the true value, and negative values where the posterior is smaller. We note that by integrating the square of this quantity over photon energy and summing over all samples we obtain the MSE that is used in the optimization. This is computed on the validation set to demonstrate generalizability. Figure 6 shows Δ_ϵ for each case 0–15 in the validation set as a function of photon energy for each configuration. The curves are offset artificially for clarity. The plot on the left shows the results for the ‘MagLIF’ configuration, the centre plot shows the results for the ‘Expert’ configuration and the plot on the right shows the results of the ‘optimum’ found using our technique. In each plot, the dashed line indicates a perfect fit while the dark coloured line shows the median of the posterior. The shaded bands show the 95 % and 68 % credible intervals.

It is immediately apparent that the credible intervals, and thus the variance of the posterior, is smaller in the optimal case than in either of the other two cases. This is most apparent at the low and high photon energy ranges where the credible bounds diverge in both the ‘MagLIF’ and ‘Expert’ configurations for most cases. For the ‘Optimum’ case we can see that the agreement between the median and the true value is excellent from about 5000–15 000 eV photon energy for all cases, consistent with the excellent agreement shown in figure 5. The posterior tends to lie somewhat below the true spectrum at lower and higher photon energies. The 68 % credible interval fully encompasses the true value (above ~ 5000 eV) and is very tightly constrained. The median of the posterior from the ‘Expert’ configuration exhibits similarly good agreement in the middle of the spectrum, somewhat worse at lower photon energies, and somewhat better at higher photon energies,




FIGURE 5. Plots showing the posterior spectrum for three different cases in the validation set inferred using each of the different instrument configurations. The top row shows the spectra inferred using the standard MagLIF configuration, the middle row using the best guess configuration, and the bottom row the optimum. The dashed black lines show the true spectrum, the solid lines show the median of the posterior, and the shaded bands show the credible intervals.

although the credible intervals are substantially larger, indicating higher variance. We can be more quantitative by computing the $\log_{10}$ of the MSE averaged over the entire validation dataset using (3.11) which is summarized in table 2. The MSE computed on the validation set is 2.13 for the ‘MagLIF’ configuration, 2.5 for the ‘Expert’ configuration and -0.22 for the ‘Optimum’, an improvement of more than two orders of magnitude. The fact that the MSE is slightly worse for the ‘Expert’ than for the ‘MagLIF’ when it clearly provides a better median fit is due in part to the large variance at the high and low photon energies, as well as the presence of a small number of cases in the validation set that produce particularly bad fits. Additionally, looking at the MSE for each case separately, we see that the difference in MSE between these two configurations is not significant relative to variance across the dataset. The improvement in MSE for the ‘Optimum’ case is significant, and shows consistent improvement across the entire validation set. This indicates that the optimization procedure used here has successfully found an improved solution to the problem of inferring the X-ray spectrum using filtered radiation detectors. We also show the peak voltages produced in each case to demonstrate that the penalty term is producing acceptable voltages. For the ‘Optimum’ case the peak recorded voltage across the validation set is 14.4 V. Inputting this value into (3.12) we find that the recorded signal is compressed by $\sim 11\%$ which is acceptable and can be corrected for with high confidence.

To compliment the MSE, we can look at how specific features of the spectrum are fit using each configuration. Figure 7 compares (a) the inferred peak intensity of the spectrum




FIGURE 6. The difference in the posterior inferred and true spectra normalized to the true spectrum, Δ_ϵ for each of the 16 validation cases considered. The dashed black line shows a value of 0, indicating perfect agreement. The solid blue line shows the mean and the light and dark shaded regions show the 68 % and 95 % credible intervals. Curves are artificially offset for clarity.

Configuration	MSE	Peak (V)	δ_I (%)	σ_I (%)	δ_E (%)	σ_E (%)	$\delta_{T_{\text{eff}}}$ (%)	$\sigma_{T_{\text{eff}}}$ (%)
MagLIF	2.13	35.3	13.5	30	6.6	16	20	67
Expert	2.5	15.6	5.3	21	4.2	14	10.5	35
Optimum	-0.22	14.4	9.3	21	2.8	6	17.2	21

TABLE 2. Summary of performance of each of the metrics discussed on the validation dataset. The first column shows the MSE defined on the validation set and the second column shows the peak voltage produced over the validation set with the given configuration. Remaining columns summarize the features computed using (4.2) and (4.3) for the peak intensity, photon energy of the peak and the continuum slope.

with the true value, (b) the inferred photon energy of the peak to the true value and (c) the inferred high energy slope of the spectrum to the truth. In each plot the inferred value is plotted on the ordinate and the true value extracted from the HFM is plotted on the abscissa. The dashed line indicates perfect agreement. The error bars indicate the 16 %–84 % credible interval, equivalent to $\pm 1\sigma$ for a Gaussian distribution. The slope is defined as an effective temperature using the relation

$$T_{\text{eff}} = \frac{\epsilon_2 - \epsilon_1}{\log(y_\epsilon(\epsilon_1)/y_\epsilon(\epsilon_2))}, \quad (4.1)$$

where we choose $\epsilon_1 = 12$ and $\epsilon_2 = 15$ keV, the results are not sensitive to this choice as long as the photon energies are above ~ 10 keV. This expression comes from the fact that the high energy portion of the spectrum is $\propto \exp(-\epsilon/kT)$. Therefore, on a log scale, the slope of this portion of the spectrum is $\propto -1/T$. Figure 7(a) shows a small improvement of the agreement between the inferred peak intensity and the truth as well as a reduction in the confidence intervals using the optimum configuration. Improvement in the peak photon energy is more pronounced, with inferred values falling much closer to the truth than in the other two cases with significantly lower scatter and credible intervals. Finally, for the slope, which should be representative of the plasma temperature, we see somewhat reduced scatter in the points about the $y = x$ line and smaller credible intervals, but with a bias towards lower inferred values. Overall, the optimum configuration is doing a better job at reproducing specific aspects of the spectrum with reduced uncertainty. For each of these metrics and each instrument configuration we compute the average bias and standard deviation defined for a quantity q as

$$\delta_q = \frac{1}{K} \sum_{j=1}^K \sqrt{\left(\frac{\langle q \rangle^j - Q^j}{Q^j} \right)^2}, \quad (4.2)$$

$$\sigma_q = \frac{1}{K} \sum_{j=1}^K \sqrt{\frac{1}{N-1} \sum_{i=1}^N \left(\frac{q_k^j - \langle q \rangle^j}{Q^j} \right)^2}, \quad (4.3)$$

where Q^j is the true value from HFM and $\langle q \rangle^j$ is the expectation value from the posterior of q for the j th training sample. We have normalized the bias and variance by the true value of Q^j so that quantities can be quoted as relative quantities. We summarize these values in table 2 as percentages. We can see that the ‘Optimum’ configuration produces better agreement, averaged over the validation set, for some of these metrics and not others. In particular, the peak photon energy is fit with lower bias δ_E and standard deviation σ_E than the other two configurations. The peak intensity is fit with a bias δ_I in between the two other configurations, but with a standard deviation σ_I similar to the expert configuration. The bias computed for the continuum slope $\delta_{T_{\text{eff}}}$ also lies between the ‘Expert’ and ‘MagLIF’ configurations, but with a substantially reduced standard deviation $\sigma_{T_{\text{eff}}}$. This demonstrates that, averaged across the entire validation dataset, the primary benefit of the optimized configuration seems to be in reducing the variance of the posterior and producing a better fit to the photon energy of the peak in the spectrum.

4.2. Comparison with model parameters

Our ultimate goal is to use the inferred spectrum and this methodology to understand the physical properties of the plasma under study. Therefore, it is desirable to compare the posterior of the model parameters temperature and ρR_ℓ as well as the total radiated




FIGURE 7. Scatter plots showing the inferred vs. true values of the peak intensity of the spectrum (a), the location of the peak in the spectrum (b) and high energy slope of the spectrum, as defined in the text (c). The true value is shown on the abscissa and the inferred value on the ordinate. Median values inferred using the standard MagLIF configuration are shown in blue, with the best guess configuration are shown in orange and the optimum configuration are shown in green. The error bars indicate the 16%–84% credible interval.

output with the values extracted from the HFM for each configuration. In order to make this comparison we use the emissivity-weighted values of temperature and ρR_ℓ . Figure 8(a) shows the median posterior temperature on the ordinate vs. the true emissivity weighted temperature on the abscissa for the ‘MagLIF’ configuration in blue, the ‘Expert’ configuration in orange, and the optimum in green. The error bars denote the 16%–84% credible interval. We see that the credible interval is reduced over the entire validation set for the optimum case. Additionally, the median clusters about the true value with less spread than in the other cases. However, there is a small bias towards lower temperature in the median inferred value from the optimal configuration. This could, however, be alleviated with a different choice for extracting the true value, or it could be a real bias incurred by using the simplified model. Figure 8(b) shows the posterior ρR_ℓ vs. the true value from the HFM. This quantity is poorly constrained with any of the instrument configurations showing no noticeable trend with the true value. The optimal configuration does reduce the credible intervals somewhat relative to the other two configurations, but the effect is minimal. These results indicate that the liner ρR_ℓ is poorly constrained with this kind of instrument. Finally, we compare the inferred total X-ray output with the true output from the HFM for each of the configurations. All three configurations perform well on this metric, indicating that this quantity is perhaps the easiest to infer, which is not surprising given the integrated nature of the instrument. The optimal and ‘best guess’ configurations do reduce the credible interval as well as the bias compared with the ‘MagLIF’ configuration, providing an overall better inference across the validation set. Using the definitions for bias and standard deviation in (4.2) and (4.3), we compute the performance of each configuration relative to the model parameters temperature, ρR_ℓ , and total radiated output on the validation dataset. These results are summarized in table 3. These metrics indicate that even though the ‘MagLIF’ case produces the minimum bias in the ρR_ℓ , the standard deviation is very large suggesting that the differences between configurations are not particularly meaningful. The ‘Expert’ configuration minimizes the bias in the inferred temperature, but the ‘Optimum’ case minimizes the standard deviation. Finally, the ‘Optimum’ case minimizes both the bias and the standard deviation of inferred total output.




FIGURE 8. Scatter plots showing the inferred vs. true values of the temperature (a), liner areal density (b) and X-ray output (c). The true value is shown on the abscissa and the inferred value on the ordinate. Median values inferred using the standard MagLIF configuration are shown in blue, with the best guess configuration are shown orange and the optimum configuration are shown orange. The error bars indicate the 16 %–84 % credible interval.

Configuration	δ_T (%)	σ_T (%)	$\delta_{\rho R}$ (%)	$\sigma_{\rho R}$ (%)	δ_Y (%)	σ_Y (%)
MagLIF	17	55	15	45	7.6	22
Expert	12.6	39	21	55	5.5	12.5
Optimum	18	22	24	34	4.4	12.2

TABLE 3. Summary of the performance of each configuration using (4.2) and (4.3) for the two model parameters temperature and ρR_ℓ as well as the total radiated output.

Figure 9 shows a corner plot illustrating the posterior distribution of the model parameters from each diagnostic configuration for the case shown in the first column of figure 5. We show the posterior distribution of the temperature, liner ρR_ℓ and total X-ray output for the ‘MagLIF’ configuration in blue, the ‘Expert’ configuration in orange, and the ‘Optimum’ configuration in green. The solid black lines indicate the true values of each quantity extracted from the HFM. The distribution of the output is similar with all configurations, showing a small bias towards lower values in the inference, although the ‘Optimum’ configuration produces a narrower distribution. The ‘MagLIF’ configuration produces a posterior distribution for the output that is slightly wider on the low side and has a heavier tail on the high side contributing to the enhanced uncertainty. The ‘Optimum’ produces the narrowest distribution for the temperature, with the ‘Expert’ configuration peaking in a similar range, but having a heavy tail at higher temperatures. This heavy tail is likely the reason for the extracted temperature appearing higher, and therefore in seeming better agreement with the truth, than the ‘Optimum’ case. In the distribution of ρR_ℓ for the ‘Expert’ configuration we can see a bimodal distribution, with a peak at low areal density and high areal density. This feature, which is persistent across the validation set contributes to the significantly increased variance in the posterior of all parameters. It is also worth noting that the correlations between the posterior temperature and ρR_ℓ are completely different in the ‘MagLIF’ case compared with the other two. All of these observations provide strong evidence that the diagnostic configuration arrived at through our optimization procedure produces a posterior that is much more tightly constrained and better behaved than the other two configurations.




FIGURE 9. Corner plot showing the posterior distribution of model parameters for one of the cases in the validation set. The plots on the diagonal show the marginalized posteriors of temperature liner areal density and total output. The off-diagonal plots show the pairwise joint distributions. The posterior obtained using the standard MagLIF configuration, best guess configuration and optimum configuration are shown in blue, orange and green, respectively.

5. Conclusions

We have demonstrated a technique for optimizing the performance of instruments to extract physical information from experiments using a combination of Bayesian inference and BO. Using Bayesian inference as the core of our optimization metric allows us to seamlessly include the effects of the myriad uncertainties that can be present regarding instrument calibrations and other configuration parameters on our ability to extract useful information. Bayesian optimization provides an efficient and flexible way to optimize over instrument parameters that can be continuous or discrete. As a pedagogical example, we optimized the configuration of a set of radiation detectors similar to what is used to diagnose MagLIF experiments at Sandia National Laboratories. We developed a parameterized physical model of the system that is used to fit synthetic data from an ensemble of HFM simulations. The reduced model represents a typical case in the physical sciences where the model used to interpret data is lacking physics that can introduce bias. We demonstrate how to handle this by choosing an optimization metric that is unambiguous with respect to the meaning of the quantities being compared. We also demonstrate how to incorporate additional constraints in the optimization that are not

captured in the inference, allowing the experimenter to enforce additional constraints or add domain specific knowledge to the optimization.

We performed this optimization using a small subset of the HFM ensemble to train the optimization and tested it against a larger subset for validation. Two reference configurations were used for comparison. We examined the posterior distributions of a variety of quantities from each of the validation cases to show that the configuration arrived at using our optimization methodology improved both the closeness of the inferences to the truth values from the HFM as well as the uncertainty of the fits. Our optimization produced a fit to the experimentally unobserved spectrum that was dramatically improved compared with both reference cases, as measured by the MSE. Additionally, the ‘Optimum’ configuration better captured the location of the peak in the spectrum and reduced the standard deviation in the high energy slope. The posterior of the model parameters temperature and ρR_ℓ as well as the inferred total X-ray output was considerably narrower and better behaved, indicating a better overall inference. The optimized configuration did show bias in some of the inferred quantities, particularly the temperature, which provides opportunity for further improvement. These results show that tradeoffs will often need to be made and that it is difficult to construct a metric that allows an instrument to perform well on all desired measurements across a wide range of possible outcomes. In the future we will explore optimization against different Q values such as the model parameters temperature and ρR_ℓ directly, or combinations of these quantities with the spectrum or spectral features.

Detailed observation of the posterior spectra offers hints at how we might refine the measurements further. Unsurprisingly, the credible intervals are largest in the high and low energy portions of the spectrum where the intensity is lowest. The addition of information that could better constrain these portions of the spectrum could dramatically improve both the bias and variance of the posterior. As an example, a simple spectrometer could measure the slope of the high energy portion of the spectrum with high confidence. The Bayesian inference methodology allows this information to be easily incorporated in a self-consistent manner. This would manifest as an additional diagnostic g_i and observation O_i added to the graph in [figure 1](#). The methodology here can therefore be extended to not only optimize the configuration of a single instrument, but a suite of instruments, allowing experimenters and facilities to understand what new measurements, and therefore investments, will provide the most constraining information for a given application. The method developed and demonstrated here is entirely general and can be applied to a wide variety of measurement and inference tasks. However, as instruments are added the dimensionality of the optimization will increase and care must be taken when choosing an optimization algorithm that can efficiently handle high-dimensional problems. It may also become necessary to explore approximations to the Bayesian inference or other algorithms that reduce the computational cost. Finally, in the future we would like to explore the use of multi-objective optimization to optimize multiple metrics simultaneously and examine the tradeoffs between them. The results here demonstrated an overall improvement in the variance of the inference but the solution produced enhanced bias in some quantities. Instead of optimizing the MSE which is the sum of the bias and variance, one could optimize both simultaneously. This approach would provide a means to examine an ensemble of solutions that effectively explore the tradeoff between metrics, allowing the analyst to choose that which satisfies the needs of the experiment.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Material	Kapton	SS	Al	Ti	V	Fe	Co	Ni	Cu	Zn	Y	Zr	Mo	Pd	Ag	Sn

TABLE 4. Filter materials along with numerical index used in the optimization.

Acknowledgements

This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Declaration of Interests. The authors report no conflict of interest.

Editor William Dorland thanks the referees for their advice in evaluating this article.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Filter materials

Table 4 shows the list of filters allowed in the optimization along with the numerical index. Filters are designated by their elemental symbol except for the two compounds Kapton and stainless steel (SS).

Appendix B. Hyperparameter tuning

In order to set reasonable values for the hyperparameters introduced by the penalty term in (3.13) and (3.14) we performed a small scan of each. First, we fixed $\alpha = 0.5$ and scanned λ over the range of $[0, 0.2]$. In figure 10 we plot the results. For each value of λ we perform the optimization and compute the peak voltage on each detector over the entire validation set. Each optimization was performed using the same random seed. The blue circles show the maximum of these voltages over the entire validation set while the orange circles show the mean over the validation set. Clearly, increasing λ decreases both metrics of the peak voltage. We settle on a value of $\lambda = 0.15$ because it produces the minimum values in this test.

To determine a suitable value of α we fixed $\lambda = 0.15$ and performed a similar scan over values of α . The results of this scan are shown in figure 11. Again, we show the peak voltage over the entire validation set and the orange circles show the mean of the peak voltage for each case in the validation set. For this scan we see a fairly shallow curve, with a minimum value somewhere near $\alpha = 0.5$. We chose to set $\alpha = 0.25$ for this work as a balance between achieving acceptable voltages and suitable quality inferences because driving voltages too low will sacrifice signal quality. While the hyper-parameter tuning scan is admittedly small, it demonstrates that the loss term functions as intended and provides a methodology to choose specific values.




FIGURE 10. Peak voltages as a function of the hyperparameter λ . Blue circles show the peak voltage found over the entire validation dataset, orange circles show the mean of the peak voltage found for each case over the dataset.




FIGURE 11. Peak voltages as a function of the hyperparameter α . Blue circles show the peak voltage found over the entire validation dataset, orange circles show the mean of the peak voltage found for each case over the dataset.

REFERENCES

- ALPER, B., DILLON, S., EDWARDS, A.W., GILL, R.D., ROBINS, R. & WILSON, D.J. 1997 The jet soft x-ray diagnostic systems. *Rev. Sci. Instrum.* **68** (1), 778–781.
- AMPLEFORD, D.J., JONES, B., JENNINGS, C.A., HANSEN, S.B., CUNEO, M.E., HARVEY-THOMPSON, A.J., ROCHAU, G.A., COVERDALE, C.A., LASPE, A.R., FLANAGAN, T.M., *et al.* 2014 Contrasting physics in wire array Z pinch sources of 1–20 keV emission on the Z facility. *Phys. Plasmas* **21** (5), 056708.
- AUTHORS, GPYOPT. 2016 GPyOpt: a Bayesian optimization framework in python. <http://github.com/SheffieldML/GPyOpt>.
- AYMAR, R., BARABASCHI, P. & SHIMOMURA, Y. 2002 The iter design. *Plasma Phys. Control. Fusion* **44** (5), 519.
- BALANDAT, M., KARRER, B., JIANG, D.R., DAULTON, S., LETHAM, B., WILSON, A.G. & BAKSHY, E. 2020 BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems 33*. MIT press.
- BISHOP, C.M. & NASRABADI, N.M. 2006 *Pattern Recognition and Machine Learning*, vol. 4. Springer.
- CHITTENDEN, J., LEBEDEV, S., JENNINGS, C., BLAND, S. & CIARDI, A. 2004 X-ray generation mechanisms in three-dimensional simulations of wire array z-pinchs. *Plasma Phys. Control. Fusion* **46** (12B), B457.

- CIARDI, A., LEBEDEV, S., FRANK, A., BLACKMAN, E., CHITTENDEN, J., JENNINGS, C., AMPLEFORD, D., BLAND, S., BOTT, S., RAPLEY, J., *et al.* 2007 The evolution of magnetic tower jets in the laboratory. *Phys. Plasmas* **14** (5), 056501.
- COVERDALE, C., JONES, B., AMPLEFORD, D., CHITTENDEN, J., JENNINGS, C., THORNHILL, J., APRUZESE, J., CLARK, R., WHITNEY, K., DASGUPTA, A., *et al.* 2010 K-shell x-ray sources at the Z accelerator. *High Energy Density Phys.* **6** (2), 143–152.
- DELGADO-APARICIO, L.F., STUTMAN, D., TRITZ, K., FINKENTHAL, M., BELL, R., HOSEA, J., KAITA, R., LEBLANC, B., ROQUEMORE, L. & WILSON, J.R. 2007 Fast electron temperature measurements using a ‘multicolor’ optical soft x-ray array. *J. Appl. Phys.* **102** (7), 073304.
- DRAKE, R.P. 2006 *High Energy Density Physics: Fundamentals, Inertial Fusion, and Experimental Astrophysics*, 1st edn. Springer.
- EPSTEIN, R., GONCHAROV, V.N., MARSHALL, F.J., BETTI, R., NORA, R., CHRISTOPHERSON, A.R., GOLOVKIN, I.E. & MACFARLANE, J.J. 2015 X-ray continuum as a measure of pressure and fuel shell mix in compressed isobaric hydrogen implosion cores. *Phys. Plasmas* **22** (2), 022707.
- FRAZIER, P.I. 2018 *Bayesian Optimization*, chap. 11, pp. 255–278. INFORMS TutORials in Operations Research. <https://pubsonline.informs.org/doi/pdf/10.1287/educ.2018.0188>.
- GOMEZ, M.R., SLUTZ, S.A., JENNINGS, C.A., AMPLEFORD, D.J., WEIS, M.R., MYERS, C.E., YAGER-ELORRIAGA, D.A., HAHN, K.D., HANSEN, S.B., HARDING, E.C., *et al.* 2020 Performance scaling in magnetized liner inertial fusion experiments. *Phys. Rev. Lett.* **125**, 155002.
- GOMEZ, M.R., SLUTZ, S.A., SEFKOW, A.B., HAHN, K.D., HANSEN, S.B., KNAPP, P.F., SCHMIT, P.F., RUIZ, C.L., SINARS, D.B., HARDING, E.C., *et al.* 2015 Demonstration of thermonuclear conditions in magnetized liner inertial fusion experiments. *Phys. Plasmas* **22** (5), 056306.
- ITER, J. 1999 Technical basis for the iter final design report, cost review and safety analysis (FDR). ITER/EDA Documentation Series, IAEA, Vienna.
- JENNINGS, C.A., CUNEO, M.E., WAISMAN, E., SINARS, D., AMPLEFORD, D., BENNETT, G., STYGAR, W. & CHITTENDEN, J. 2010 Simulations of the implosion and stagnation of compact wire arrays. *Phys. Plasmas* **17** (9), 092703.
- JONES, D.R., SCHONLAU, M. & WELCH, W.J. 1998 Efficient global optimization of expensive black-box functions. *J. Global Optim.* **13** (4), 455–492.
- JONES, M., AMPLEFORD, D., CUNEO, M., HOHLFELDER, R., JENNINGS, C., JOHNSON, D., JONES, B., LOPEZ, M., MACARTHUR, J., MILLS, J., *et al.* 2014 X-ray power and yield measurements at the refurbished Z machine. *Rev. Sci. Instrum.* **85** (8), 083501.
- JOSEPH, V.R. & VAKAYIL, A. 2022 Split: an optimal method for data splitting. *Technometrics* **64** (2), 166–176.
- KEILHACKER, M. & JET TEAM 1999 Fusion physics progress on the joint European torus (jet). *Plasma Phys. Control. Fusion* **41** (12B), B1.
- KNAPP, P., GLINSKY, M., SCHAEUBLE, M., JENNINGS, C., EVANS, M., GUNNING, J., AWE, T., CHANDLER, G., GEISSEL, M., GOMEZ, M., *et al.* 2022 Estimation of stagnation performance metrics in magnetized liner inertial fusion experiments using Bayesian data assimilation. *Phys. Plasmas* **29** (5), 052711.
- KNAPP, P., GOMEZ, M., HANSEN, S., GLINSKY, M., JENNINGS, C., SLUTZ, S., HARDING, E., HAHN, K., WEIS, M., EVANS, M., *et al.* 2019 Origins and effects of mix on magnetized liner inertial fusion target performance. *Phys. Plasmas* **26** (1), 012704.
- KNUDDE, N., VAN DER HERTEN, J., DHAENE, T. & COUCKUYT, I. 2017 Gpflowopt: a Bayesian optimization library using tensorflow.
- KOCHENDERFER, M.J. & WHEELER, T.A. 2019 *Algorithms for Optimization*. Mit Press.
- LINDL, J. 1995 Development of the indirect – drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Phys. Plasmas* **2** (11), 3933–4024.
- LUXON, J.L. 2002 A design retrospective of the diii-d tokamak. *Nucl. Fusion* **42** (5), 614.
- MAK, S. & JOSEPH, V.R. 2018 Support points. *Ann. Statist.* **46** (6A), 2562–2592.
- MA, T., PATEL, P.K., IZUMI, N., SPRINGER, P.T., KEY, M.H., ATHERTON, L.J., BENEDETTI, L.R., BRADLEY, D.K., CALLAHAN, D.A., CELLIERS, P.M., *et al.* 2013 Onset of hydrodynamic mix in high-velocity, highly compressed inertial confinement fusion implosions. *Phys. Rev. Lett.* **111**, 085004.

- MATZEN, M.K., SWEENEY, M.A., ADAMS, R.G., ASAY, J.R., BAILEY, J.E., BENNETT, G.R., BLISS, D.E., BLOOMQUIST, D.D., BRUNNER, T.A., CAMPBELL, R.B., *et al.* 2005 Pulsed-power-driven high energy density physics and inertial confinement fusion research. *Phys. Plasmas* **12** (5), 055503.
- MCKAY, M.D., CONOVER, W.J. & BECKMAN, R.J. 1979 A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**, 239–245.
- MITCHELL, M. 1998 *An Introduction to Genetic Algorithms*. MIT Press.
- MURPHY, K.P. 2022 *Probabilistic Machine Learning: An Introduction*. MIT Press.
- NUCKOLLS, J., WOOD, L., THIESSEN, A. & ZIMMERMAN, G. 1972 Laser compression of matter to super-high densities: thermonuclear (CTR) applications. *Nature* **239** (5368), 139–142.
- PATIL, A., HUARD, D. & FONNESBECK, C.J. 2010 Pymc: Bayesian stochastic modelling in python. *J. Stat. Softw.* **35** (4), 1.
- REBUT, P., BICKERTON, R. & KEEN, B.E. 1985 The joint European torus: installation, first results and prospects. *Nucl. Fusion* **25** (9), 1011.
- SAVAGE, M.E., LECHIEN, K.R., LOPEZ, M.R., STOLTZFUS, B.S., STYGAR, W.A., ARTERY, D.S., LOTT, J.A. & CORCORAN, P.A. 2011 Status of the Z pulsed power driver. In *2011 IEEE Pulsed Power Conference*, pp. 983–990.
- SCHULZ, E., SPEEKENBRINK, M. & KRAUSE, A. 2018 A tutorial on Gaussian process regression: modelling, exploring, and exploiting functions. *J. Math. Psychol.* **85**, 1–16.
- SHAHRIARI, B., SWERSKY, K., WANG, Z., ADAMS, R.P. & DE FREITAS, N. 2016 Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104** (1), 148–175.
- SILVEY, S. 1980 *Optimal Design: An Introduction to the Theory for Parameter Estimation*. Chapman & Hall.
- SLUTZ, S.A., HERRMANN, M.C., VESEY, R.A., SEFKOW, A.B., SINARS, D.B., ROVANG, D.C., PETERSON, K.J. & CUNEO, M.E. 2010 Pulsed-power-driven cylindrical liner implosions of laser preheated fuel magnetized with an axial field. *Phys. Plasmas* **17** (5), 056303.
- SPIELMAN, R.B., HSING, W.W. & HANSON, D.L. 1988 Photoconducting x-ray detectors for z-pinch experiments. *Rev. Sci. Instrum.* **59** (8), 1804–1806.
- VON TOUSSAINT, U. 2011 Bayesian inference in physics. *Rev. Mod. Phys.* **83** (3), 943.
- WIKLE, C.K. & BERLINER, L.M. 2007 A Bayesian tutorial for data assimilation. *Physica D* **230** (1–2), 1–16.
- WILLIAMS, C.K. & RASMUSSEN, C.E. 2006 *Gaussian Processes for Machine Learning*, vol. 2. MIT Press.
- YAGER-ELORRIAGA, D., GOMEZ, M.R., RUIZ, D.E., SLUTZ, S.A., HARVEY-THOMPSON, A.J., JENNINGS, C., KNAPP, P., SCHMIT, P., WEIS, M., AWE, T.J., *et al.* 2021 An overview of magneto-inertial fusion on the Z machine at sandia nationallaboratories. *Nucl. Fusion* **62** (4).