Review of Scientific Instruments ARTICLE pubs.aip.org/aip/rsi

Quantification and visualization of uncertainties in reconstructed penumbral images of implosions at Omega

Cite as: Rev. Sci. Instrum. 95, 063506 (2024); doi: 10.1063/5.0214641
Submitted: 18 April 2024 • Accepted: 8 June 2024 •
Published Online: 24 June 2024

View Online Export Citation CrossMark

J. H. Kunimune,[1,a)] P. V. Heuer,[2] B. L. Reichelt,[1] T. M. Johnson,[1] and J. A. Frenje[1]

AFFILIATIONS

[1] Plasma Science and Fusion Center, Massachusetts Institute of Technology, 167 Albany St., Cambridge, Massachusetts 02139, USA

[2] Laboratory for Laser Energetics, University of Rochester, 250 E. River Rd., Rochester, New York 14623, USA

Note: This paper is part of the Special Topic on Proceedings of the 25th Topical Conference on High-Temperature Plasma Diagnostics.

[a)] **Author to whom correspondence should be addressed:** kunimune@mit.edu

ABSTRACT

Penumbral imaging is a technique used in plasma diagnostics in which a radiation source shines through one or more large apertures onto a detector. To interpret a penumbral image, one must reconstruct it to recover the original source. The inferred source always has some error due to noise in the image and uncertainty in the instrument geometry. Interpreting the inferred source thus requires quantification of that inference's uncertainty. Markov chain Monte Carlo algorithms have been used to quantify uncertainty for similar problems but have never been used for the inference of the shape of an image. Because of this, there are no commonly accepted ways of visualizing uncertainty in two-dimensional data. This paper demonstrates the application of the Hamiltonian Monte Carlo algorithm to the reconstruction of penumbral images of fusion implosions and presents ways to visualize the uncertainty in the reconstructed source. This methodology enables more rigorous analysis of penumbral images than has been done in the past.

© 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/). https://doi.org/10.1063/5.0214641

24 June 2024 15:10:59

I. INTRODUCTION

Penumbral imaging is a form of coded-aperture imaging that is commonly used in plasma diagnostics. A radiation source with an unknown size and shape, such as an inertially confined fusion (ICF) implosion, projects radiation through one or more circular apertures onto a detector, such as an image plate, which records an image of the source. In the field of ICF, penumbral imaging is used to image emission of protons,[1] deuterons,[2,3] neutrons,[4] x rays,[5,6] and gamma-rays.[7]

Penumbral imaging is distinguished from pinhole imaging in that the apertures are larger than the source. This increases the quantity of radiation measured and thus improves the quality of the image but also makes the image more difficult to interpret. Unlike a pinhole image, it is not a simple picture of the source. Instead, it is a *penumbral image*—a convolution of the source with a point-spread function—that must be reconstructed. An example of a penumbral image is shown in Fig. 1. It looks like an image of the aperture, and information about the source is encoded in the circular edge. There are multiple ways to reconstruct a penumbral image, including Fourier transform based methods, such as Wiener deconvolution,[5] computed tomography methods, such as filtered backprojection,[1,6] and iterative maximum-likelihood methods, such as the Richardson–Lucy algorithm.[8,9]

The recorded penumbral image will always have errors due to random noise in the image and uncertainty in the point-spread function, of which both propagate into the inferred source. There are ways to reduce the random error, but they all require tuning to discriminate noise from real features in the data. For example, Richardson–Lucy can be initialized with a uniform source and terminated after a finite number of iterations, but the user must decide which iteration to terminate at.[3] With too much noise reduction, some features of the true source may be lost in the inferred source; with too little noise reduction, some features may appear in the

© Author(s) 2024

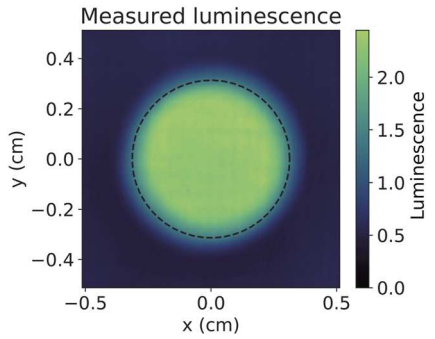

FIG. 1. A penumbral x-ray image of an ICF implosion experiment at Omega.

inferred source that were not present in the true source. This is illustrated in Fig. 2.

Monte Carlo methods are often used to quantify the uncertainty in these reconstructions. Two such methods are bootstrapping[10] and adding noise to the image.[11] In both, a set of penumbral images is generated, all based on the data but with some variation, and each is reconstructed individually. Variations in the reconstructed sources quantify the error in the inference due to noise in the image and due to random variations in the aperture geometries. However, neither method accounts for systematic errors due to the reconstruction algorithm.

An alternative method is the application of Bayesian statistics, in which the true value of the parameter being inferred is treated as a random variable with a probability distribution. In the context of penumbral imaging, this means that rather than varying the image and seeing how it affects the inferred source, the image is kept constant and the source is varied in such a way that it remains consistent with the image. This framework allows random and systematic uncertainties to be quantified and assigned probability distributions.

In problems with many degrees of freedom such as image reconstruction, Bayesian statistics is usually implemented through Markov chain Monte Carlo methods, which allow a many-dimensional probability distribution to be sampled. Such methods have been used with "L rolled edge" penumbral imaging to quantify the size uncertainty of an x-ray backlighter.[12] However, they are rarely used for imaging problems in general and have never been used to quantify uncertainty in the *shape* of a reconstructed source. One question that arises is how to visualize the resulting distribution of sources, as traditional error bars are not possible in two-dimensional data.

This paper demonstrates the use of Markov chain Monte Carlo methods to quantify uncertainty in the reconstruction of penumbral images of ICF implosions. Various ways of visualizing the uncertainty in the source are presented.

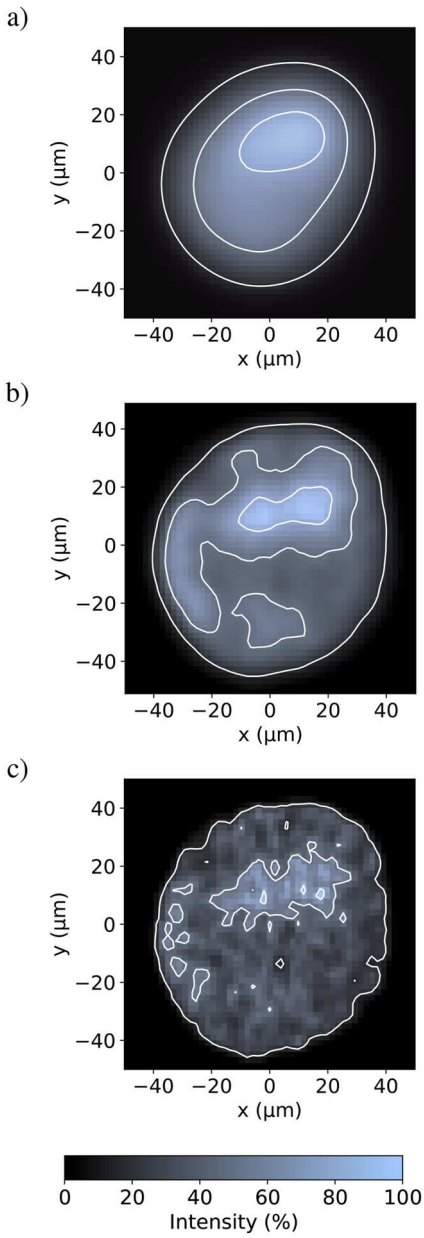

FIG. 2. Effect of the number of Richardson–Lucy iterations on the level of detail in an inferred source. Shown is the source reconstructed from the image in Fig. 1 with (a) 30, (b) 370, or (c) 3000 iterations. Isocontours are shown for 17%, 50%, and 83% of the peak intensity.

II. UNCERTAINTY QUANTIFICATION

In Bayesian statistics, Bayes' theorem is used to calculate the posterior probability distribution of the set of unobserved source pixels $\{s_{ij}\} = \mathbf{s}$ given a set of observed image pixels $\{d_{kl}\} = \mathbf{d}$,

$$p(\mathbf{s}|\mathbf{d}) = C_P \cdot p(\mathbf{s}) \cdot p(\mathbf{d}|\mathbf{s}). \quad (1)$$

Here, C_P is a normalization constant set such that the probability distribution integrates to 1. For most analyses, it need not be calculated. All that is needed are methods to calculate the prior probability $p(\mathbf{s})$ and the likelihood $p(\mathbf{d}|\mathbf{s})$ and to sample the resulting distribution $p(\mathbf{s}|\mathbf{d})$.

24 June 2024 15:10:59

© Author(s) 2024

A. Likelihood

First, to quantify $p(\mathbf{d}|\mathbf{s})$, the relationship between the image and source is expressed linearly as

$$\mathbf{d} = \mathbf{A}\mathbf{s} + d_{\mathrm{B}} + \boldsymbol{\varepsilon}. \tag{2}$$

Here, $\mathbf{A}$ represents convolution with the point-spread function, d_{B} is the image background, and $\boldsymbol{\varepsilon}$ is the pixel noise. This does not account for the spatial resolution of the image plate.

The background arises primarily from neutrons and gamma-rays from the implosion that interact with the detector. Since, unlike the x rays being imaged, they do not interact with the aperture array, their fluence on the detector is uniform, and their effect on the image can be represented by a single scalar.

For imaging of ICF implosions, the noise values ε_{kl} are dominated by photon counting statistics, so a Poisson likelihood function would be appropriate. In practice, it is difficult to calculate the absolute number of photons detected in each pixel, so this is challenging to implement. Instead, the Poisson distribution is approximated by a Gaussian distribution whose standard deviation σ_{kl} scales with the square root of the pixel value. The magnitude of the noise is estimated from the standard deviation of pixel values in the flat center of the image, which is expressed mathematically as

$$\sigma_{kl} = \sqrt{\frac{\left\langle (\mathbf{d} - \langle \mathbf{d} \rangle)^2 \right\rangle}{\langle \mathbf{d} \rangle} d_{kl}}. \tag{3}$$

Here, the angle brackets represent an average over the flat center of the image.

This enables the calculation of the likelihood of the data,

$$p(\mathbf{d}|\mathbf{s}) = C_d \exp\left(-\frac{1}{2} \sum_{k,l} \left(\frac{(\mathbf{d} - \mathbf{A}\mathbf{s})_{kl} - d_{\mathrm{B}}}{\sigma_{kl}} \right)^2 \right). \tag{4}$$

B. Prior

Next, the prior probability must be defined for the source, as well as for any other parameters that affect the likelihood. For the background d_{B}, an uninformative gamma distribution is chosen. This distribution forbids negative values and has a long tail that permits very large values,

$$p(d_{\mathrm{B}}) = \begin{cases} 0 & \text{if } d_{\mathrm{B}} \le 0, \\ C_{\mathrm{B}}\, d_{\mathrm{B}}^{-1/2} \exp\left(-d_{\mathrm{B}}/\langle d \rangle\right) & \text{if } d_{\mathrm{B}} > 0. \end{cases} \tag{5}$$

For the source, the simplest prior would treat each pixel as an independent random variable. However, this results in predictions with large pixel-to-pixel variations, which are unrealistic. Since the true source is a physical object, it stands to reason that there should be some correlation between nearby pixels. The simplest way to enforce such a correlation is to penalize strong gradients according to some smoothing parameter α, as proposed by Fowler *et al.*,[12]

$$p(\mathbf{s}|\alpha) = C_s \exp\left(-\alpha \sum_{i,j} |\nabla s_{ij}|^2 \right), \tag{6}$$

$$|\nabla s_{ij}|^2 = \left(\frac{s_{i+1,j} - s_{i,j}}{\Delta x} \right)^2 + \left(\frac{s_{i,j+1} - s_{i,j}}{\Delta y} \right)^2. \tag{7}$$

All source pixel values are constrained to be non-negative, to ensure that the results are physically possible. All pixels along the edge of the source domain are also fixed to zero. This ensures that the source prior's covariance matrix is nonsingular. These constraints can be imposed mathematically by setting the prior to zero anywhere they are violated, but, computationally, it is easier to impose them through the parameterization of the source (for example, by sampling the logarithms of the pixel values rather than the pixel values themselves).

The smoothing parameter α could be set manually, but this would result in the same unquantified systematic uncertainty that is present in existing reconstruction algorithms. Instead, α is treated as an unknown variable with its own prior probability distribution. A ballpark estimate is obtained by estimating the mean value of $\mathbf{s}$ from the intensity in the flat center of the image and noting that, by dimensional analysis, α should scale inversely with it. A wide log-Gaussian distribution centered on that estimate is then chosen to permit a wide range of αs over multiple orders of magnitude,

$$p(\alpha) = \frac{C_\alpha}{\alpha} \exp\left(-\frac{1}{2} \left(\frac{\ln \alpha - \ln \alpha_0}{3} \right)^2 \right), \tag{8}$$

$$\alpha_0 = \frac{\Delta x \Delta y}{\left(\frac{\langle d \rangle}{n^2} \right)^2}. \tag{9}$$

Here, n^2 is the number of free pixels in the source.

In this scheme, the values of α that are too small are penalized because the solution space is so big and only a small part of it fits the data. The values of α that are too big are penalized because it cannot fit the data without incurring large gradient penalties. Thus, the posterior distribution will naturally cluster around the value of α that reduces the solution space as much as possible without excluding the solutions that fit the data.

Note that the constant C_s cannot be ignored because it varies with α, so its dependence on α must be quantified. Since $p(\mathbf{s})$ is an n^2-dimensional Gaussian distribution where every element of the covariance matrix is inversely proportional to α, it can be shown that $C_s \propto \alpha^{n^2/2}$. This factor is used in the calculation of $p(\mathbf{s}|\alpha)$ in place of C_s.

Equation (1) must be modified to allow variation in the parameters d_{B} and α. The full joint posterior probability distribution is calculated as follows:

$$p(\mathbf{s}, \alpha, d_{\mathrm{B}}|\mathbf{d}) = C_{\mathrm{J}}\, p(\mathbf{s}|\alpha)\, p(\alpha)\, p(d_{\mathrm{B}})\, p(\mathbf{d}|\mathbf{s}, d_{\mathrm{B}}). \tag{10}$$

The point-spread function $\mathbf{A}$ can be parameterized in terms of the aperture radius and shape and thus also be treated as a random variable and assigned *a prior*. This would enable the quantification of systematic error due to the uncertainty in the point-spread function and is the subject of future work.

C. Sampling

For the starting point of the Markov chain, the parameters α and d_{B} are set to the medians of their respective prior distributions.

24 June 2024 15:10:59

© Author(s) 2024

The free pixels of the source are initialized by drawing from a log-uniform distribution between $\frac{\langle d \rangle}{n^2} e^{-1}$ and $\frac{\langle d \rangle}{n^2} e$. This ensures that the sum of the source pixels is approximately correct while providing no shape information.

The Markov chain is executed using PyMC,[13] an open-source Bayesian statistics library for Python. PyMC implements the no-U-turn sampler,[14] a self-tuning implementation of the Hamiltonian Monte Carlo algorithm.[15] It automatically calculates and differentiates the unnormalized posterior probability density and samples from the posterior distribution.

Eight Markov chains are run in parallel for 3000 iterations each, consuming a total of 5.5 CPU-h. The first 2000 iterations are used to tune the step size and are discarded after the sampling is complete. That leaves a total of 1000 posterior samples in each Markov chain.

III. UNCERTAINTY VISUALIZATION

Here, an example is shown of this methodology applied to hard x-ray data collected from the Omega Laser Facility.[16] The source was a laser direct-drive ICF implosion. The target was a DT gas-filled spherical shell of deuterated plastic with a diameter of 870 μm. It was offset 100 μm from the target chamber center, roughly in the direction of the imager, to induce an asymmetry. The detector was an image plate filtered with 10 μm of tantalum and 50 μm of aluminum, which captured seven nominally identical projected images from apertures 300 μm in diameter. These were aligned by fitting their centroids to the known shape of the aperture array, resampled to a common pixel grid, and summed. The result was analyzed as a single 259×259-pixel image. The resulting source has 97×97 free pixels, each 2×2 μm^2 in size.

The simplest way to interpret the results is to examine the individual sources in the Markov chains. Figure 3 shows nine such randomly chosen sources. This provides an intuitive sense of the significant features and plausible variations in the source. For instance, it can be noted that all nine sources exhibit a strong skewness, with the intensity peak up and to the right of the centroid of the 17% contour.

For most practical interpretations, it is preferable to examine a single representative source and the magnitude of its plausible error. For zero- or one-dimensional data, this can be done using error bars or credible bands. This becomes more difficult for two-dimensional images.

This can be addressed by using one-dimensional representations of the image, such as lineouts. Figure 4 shows a lineout through the normalized source along the line $x = 0$. Specifically, a pointwise equal-tailed 90% credible band is shown, which characterizes the range of plausible locations of the lineout. At each x-value, the distribution of possible y-values is calculated, and shading is applied between that distribution's fifth percentile and 95th percentile. In other words, shading is applied to every point on the plot that is neither 95% likely to be below the true lineout, nor 95% likely to be above the true lineout. A lineout of a single representative sample from the Markov chains is also shown.

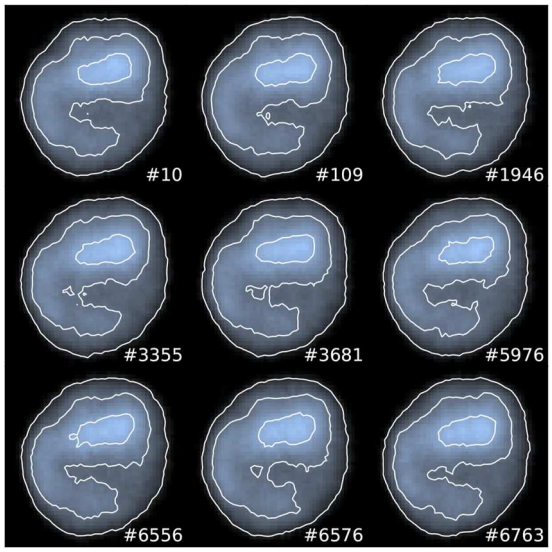

FIG. 3. Nine samples randomly selected from the 8000 iterations produced by the Markov chains.

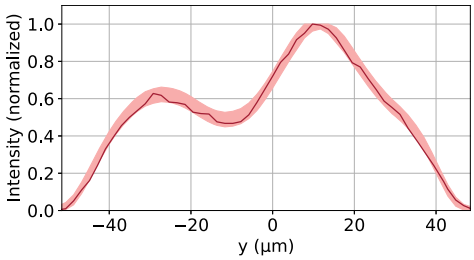

FIG. 4. In red, a vertical source lineout taken from a single representative sample. In pink, its pointwise equal-tailed 90% credible band.

Another popular one-dimensional representation is a contour. In particular, the contour at 17% of the peak intensity is a standard way to characterize images of ICF implosions. A pointwise equal-tailed 90% credible band can be constructed for a contour in a manner analogous to the lineout. Shading is applied to every point that is neither 95% likely to be within the true contour nor 95% likely to be outside the true contour. The resulting shaded region characterizes the range of plausible locations of the contour. Figure 5 shows credible bands for the 17%-, 50%-, and 83%-contours of the normalized source.

A more direct way to visualize the posterior distribution is to extract important scalars and present their possible values as a histogram. Figure 6 shows the result of this applied to the source's 17% contour's mean radius and its relative mode-2 asymmetry, of which both are commonly used to quantitatively describe images of ICF implosions.

24 June 2024 15:10:59

© Author(s) 2024

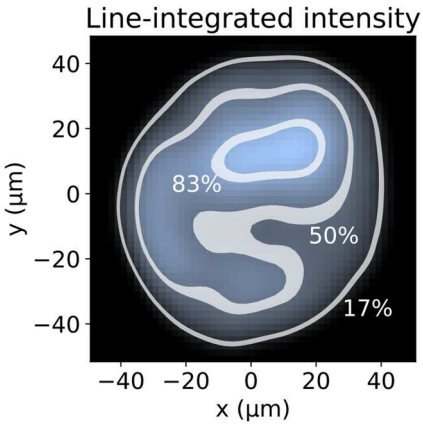

FIG. 5. Pixelwise median of the distribution of normalized sources, with pointwise 90% credible bands shown for the contours at 17%, 50%, and 83% of the peak intensity.

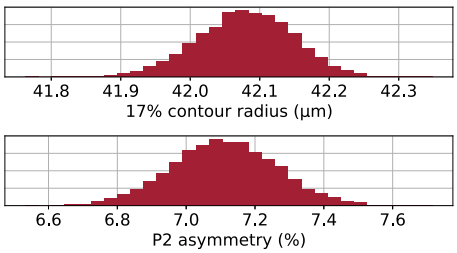

FIG. 6. Posterior histograms of two scalars that characterize the source size and shape. The contour radius is 90% likely to be between 41.90 and 42.23 μm, and the asymmetry is 90% likely to be between 6.84% and 7.42%.

IV. CONCLUSIONS

Penumbral imaging is a technique often used in plasma diagnostics in which a radiation source projects radiation through one or more apertures onto a detector, where the apertures are larger than the source.[5] To interpret a penumbral image, one must first use a reconstruction algorithm to recover the original source. The inferred source always has some error due to noise in the image and uncertainty in the point-spread function. There are methods to reduce this random error, but they are challenging to tune and they introduce their own systematic errors. Interpreting the inferred source thus requires quantification of that inference's uncertainty. However, existing uncertainty quantification methods do not account for systematic error from the algorithm.[10,11]

Bayesian Markov chain Monte Carlo algorithms such as Hamiltonian Monte Carlo[15] have been used to quantify the uncertainty in reconstructed penumbral images of x-ray backlighter sources,[12] but have never been used for penumbral images of ICF implosions. In particular, previous applications of Markov chain Monte Carlo have only quantified uncertainty in source sizes, not in source shapes. As a result, there are no common means of visualizing uncertainty in image data.

This paper demonstrates the application of Hamiltonian Monte Carlo to the reconstruction of penumbral images of ICF implosions. Multiple methods to visualize the uncertainty in the reconstructed source are presented, including credible bands of lineouts and contours and histograms of important scalars. This methodology allows for more rigorous quantitative interpretation of all coded-aperture imaging.

ACKNOWLEDGMENTS

This work was supported, in part, by the Laboratory for Laser Energetics under Contract No. 417532G/UR FAO GR510907, the U.S. Department of Energy NNSA MIT Center-of-Excellence under Contract No. DE-NA0003868, and the NNSA Laboratory Residency Graduate Fellowship under Contract No. DE-NA0003960.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

J. H. Kunimune: Conceptualization (equal); Investigation (lead). **P. V. Heuer**: Investigation (supporting). **B. L. Reichelt**: Conceptualization (equal). **T. M. Johnson**: Resources (lead). **J. A. Frenje**: Supervision (lead).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1]F. H. Séguin, J. L. DeCiantis, J. A. Frenje, S. Kurebayashi, C. K. Li, J. R. Rygg, C. Chen, V. Berube, B. E. Schwartz, R. D. Petrasso, V. A. Smalyuk, F. J. Marshall, J. P. Knauer, J. A. Delettrez, P. W. McKenty, D. D. Meyerhofer, S. Roberts, T. C. Sangster, K. Mikaelian, and H. S. Park, Rev. Sci. Instrum. **75**, 3520 (2004).
[2]J. H. Kunimune, H. G. Rinderknecht, P. J. Adrian, P. V. Heuer, S. P. Regan, F. H. Séguin, M. Gatu Johnson, R. P. Bahukutumbi, J. P. Knauer, B. L. Bachmann, and J. A. Frenje, Phys. Plasmas **29**, 072711 (2022).
[3]H. G. Rinderknecht, P. V. Heuer, J. Kunimune, P. J. Adrian, J. P. Knauer, W. Theobald, R. Fairbanks, B. Brannon, L. Ceurvorst, V. Gopalaswamy, C. A. Williams, P. B. Radha, S. P. Regan, M. G. Johnson, F. H. Séguin, and J. A. Frenje, Rev. Sci. Instrum. **93**, 093507 (2022).
[4]F. E. Merrill, D. Bower, R. Buckles, D. D. Clark, C. R. Danly, O. B. Drury, J. M. Dzenitis, V. E. Fatherley, D. N. Fittinghoff, R. Gallegos, G. P. Grim, N. Guler, E. N. Loomis, S. Lutz, R. M. Malone, D. D. Martinson, D. Mares, D. J. Morley, G. L. Morgan, J. A. Oertel, I. L. Tregillis, P. L. Volegov, P. B. Weiss, C. H. Wilde, and D. C. Wilson, Rev. Sci. Instrum. **83**, 10D317 (2012).
[5]P. J. Adrian, B. Bachmann, R. Betti, A. Birkel, P. Heuer, M. G. Johnson, N. V. Kabadi, J. P. Knauer, J. Kunimune, C. K. Li, O. M. Mannion, R. D. Petrasso, S. P. Regan, H. G. Rinderknecht, C. Stoeckl, F. Séguin, A. Sorce, R. C. Shah, G. D. Sutcliffe, and J. A. Frenje, Rev. Sci. Instrum. **93**, 113540 (2022).
[6]B. Bachmann, T. Hilsabeck, J. Field, N. Masters, C. Reed, T. Pardini, J. R. Rygg, N. Alexander, L. R. Benedetti, T. Döppner, A. Forsman, N. Izumi, S. LePape, T. Ma, A. G. MacPhee, S. R. Nagel, P. Patel, B. Spears, and O. L. Landen, Rev. Sci. Instrum. **87**, 11E201 (2016).
[7]V. E. Fatherley, D. N. Fittinghoff, R. L. Hibbard, H. J. Jorgenson, J. I. Martinez, J. A. Oertel, D. W. Schmidt, C. S. Waltz, C. H. Wilde, and P. L. Volegov, Rev. Sci. Instrum. **89**, 10I127 (2018).

24 June 2024 15:10:59

© Author(s) 2024

[8]W. H. Richardson, J. Opt. Soc. Am. **62**, 55 (1972).
[9]V. I. Gelfgat, E. L. Kosarev, and E. R. Podolyak, Comput. Phys. Commun. **74**, 335 (1993).
[10]K. M. Lamb, V. Geppert-Kleinrath, N. W. Birge, C. R. Danly, L. Divol, D. N. Fittinghoff, M. S. Freeman, A. E. Pak, C. H. Wilde, A. B. Zylstra, and P. L. Volegov, Rev. Sci. Instrum. **93**, 043508 (2022).
[11]K. W. Wong and B. Bachmann, Rev. Sci. Instrum. **93**, 073501 (2022).
[12]M. J. Fowler, M. Howard, A. Luttman, S. E. Mitchell, and T. J. Webb, Inverse Probl. Sci. Eng. **24**, 353–371 (2016).
[13]J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, PeerJ Comput. Sci. **2**, e55 (2016).
[14]M. D. Hoffman and A. Gelman, J. Mach. Learn. Res. **15**, 1593 (2014).
[15]S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, Phys. Lett. B **195**, 216 (1987).
[16]T. Boehly, D. Brown, R. Craxton, R. Keck, J. Knauer, J. Kelly, T. Kessler, S. Kumpan, S. Loucks, S. Letzring, F. Marshall, R. McCrory, S. Morse, W. Seka, J. Soures, and C. Verdon, Opt. Commun. **133**, 495 (1997).

24 June 2024 15:10:59

© Author(s) 2024