

Description of complex viewing geometries of fusion tomography diagnostics by ray-tracing

M. Carr,^{1,a)} A. Meakins,¹ M. Bernert,² P. David,² C. Giroud,¹ J. Harrison,¹ S. Henderson,¹ B. Lipschultz,³ F. Reimold,⁴ EUROfusion MST1 Team,^{b)} and ASDEX Upgrade Team^{c)}

¹CCFE, Culham Science Centre, Abingdon, Oxon OX14 3DB, United Kingdom

²Max-Planck-Institut für Plasmaphysik, 85748 Garching, Germany

³Department of Physics, York Plasma Institute, University of York, Heslington, York, United Kingdom

⁴Max-Planck-Institut für Plasmaphysik, Greifswald, Germany

(Received 28 March 2018; accepted 24 July 2018; published online 16 August 2018)

Ray-tracing techniques are applied to bolometry, a diagnostic where the finite collection volume is particularly sensitive to the machine and detector configuration. A technique is presented that can handle arbitrarily complex aperture and collimator geometries, neglecting reflection effects. Sight lines from the ASDEX Upgrade bolometer foils were ray-traced with a path tracing algorithm, where the optical path is represented by a statistical bundle of ray paths connecting the foil surface with the slit geometry. By using the full 3D machine model for the detector box and first wall, effects such as occlusion and vignetting were included in the calculation of the bolometer's étendue. Inversion matrices calculated with the ray-tracing technique were compared with the more conventional single-ray approach and shown to be naturally more constrained, requiring less regularisation. The two models were tested on a sample radiation scenario, and the common single-ray approximation is shown to be insufficient. These results are particularly relevant for the divertor where strong emission gradients may be present. The technique developed generalises well to arbitrarily complex viewing geometries and collimators, opening up a new design space for bolometer configurations that might not normally have been considered. <https://doi.org/10.1063/1.5031087>

I. INTRODUCTION

To protect the divertor components in machines on the scale of ITER and DEMO, a large fraction of the exhaust power entering the scrape off layer must be radiated before it reaches the divertor targets.¹ Multi-channel bolometry is a key diagnostic used in current experiments to facilitate the development of highly radiating divertor scenarios.² An accurate bolometric diagnostic is essential for both measuring the total radiated power fraction and determining the spatial distribution of the emissivity.

A bolometer is designed to effectively be a black surface such that it will absorb all photons incident upon it, yielding a measurement of the radiant flux at that surface. Gold is usually used as the absorbing material of choice because it exhibits strong absorption above ~ 3 eV, which corresponds well to the spectral region where the bulk of the power is radiated in a typical tokamak plasma.²

The limited geometric accessibility around tokamak plasmas means that most bolometry diagnostics are installed as fans of pinhole cameras at a low number of observation points. The plasma's local emissivity profile is then inferred from the line integrated measurements by means of standard tomography techniques.^{2,3} Figure 1 shows the FLH bolometer camera *in situ* at ASDEX Upgrade (AUG).

The response of the collimated detectors is often approximated as infinitely thin straight lines, which will hereafter be referred to as the single-ray model. However, this approximation is known to be a poor representation of the detector when the beam width is significantly larger than the length scale of the emission features.⁴⁻⁸

The most advanced analytical model was developed by Ingesson *et al.*⁴ and considers the geometrical response of an ideal rectangular pinhole and detector. The authors also developed a numerical treatment for calculating the response of the as-installed JET bolometer detectors which consisted of a rectangular detector with a cylindrical collimator. Occlusion of the detection volume by first wall protection tiles was taken into account through secondary apertures.⁴ In later work, it was shown how this model could be used to optimise the design of apertures and collimators for two dimensional tomography.⁶ However, this model was not able to take into account the toroidal bending of the voxels due to the nature of the numerical schemes used. Although negligible for JET, this effect has been shown to be significant on other devices.⁷ Another numerical scheme calculates the sensitivity volume through a 3d integration grid consisting of lattice planes aligned parallel to the detector surface.^{7,8} At each integration point, the solid angle of the detector is calculated including partial shadowing by the pinhole.

The most complete numerical approach would be to use Monte Carlo techniques to launch an ensemble of rays from the detector surface and trace them through to collisions with the as-built engineering model for detector aperture surfaces and the first wall. This process would result in a complete model of the direct coupling between the emission source and

^{a)}matthew.carr@ukaea.uk

^{b)}See the author list in H. Meyer *et al.* "Overview of progress in European medium sized tokamaks towards an integrated plasma-edge/wall solution," Nucl. Fusion 57, 102014 (2017).

^{c)}See the author list in A. Kallenbach *et al.* "Overview of ASDEX Upgrade results," Nucl. Fusion 57, 102015 (2017).

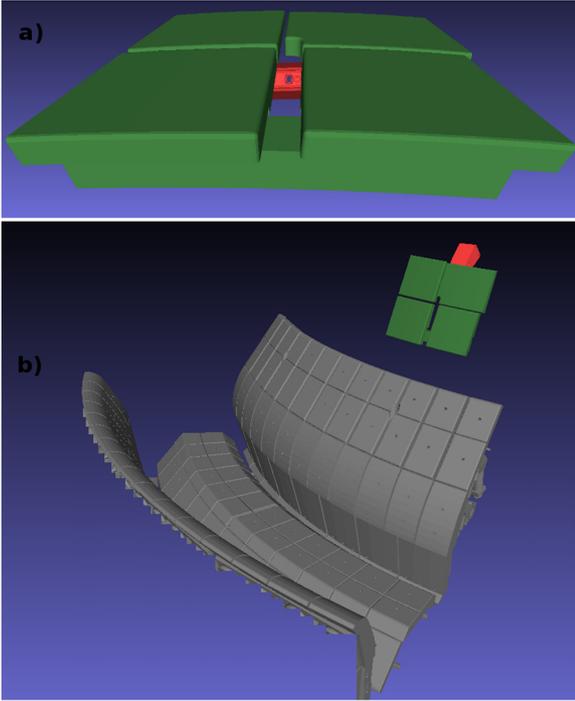


FIG. 1. (a) Example visualisation of the FLH bolometer camera (red), which contains four bolometer foil sight lines, with its pinhole viewing between a gap in the first wall tiles (green). In (b), the same FLH camera is shown with its context viewing into the AUG divertor (gray). For scale, the red FLH bolometer camera is 10 cm long, and the gap between the tiles is approximately 1.8 cm. Note that the model detail has been reduced for visualisation.

the detector including all 3D effects, such as occlusion and vignetting: Occlusion being when rays are obstructed by structures external to the camera (e.g., wall tiles) and vignetting being rays obstructed by the camera aperture. This Monte Carlo approach will be referred to as the volume ray-tracing technique.

The CHERAB code^{9,10} was developed at JET as a platform for modeling spectroscopic diagnostics with the Raysect ray-tracing package.¹¹ Raysect is a fully spectral open-source scientific ray-tracing framework that is capable of handling detailed 3D engineering geometries and physically accurate reflections. In this work, CHERAB has been extended for use with bolometry by exploiting Raysect's Monte Carlo ray-tracing capabilities. Although Raysect is capable of modeling reflection effects as well, these models require detailed measurements of the reflecting properties of many in-vessel components. The impact of reflections will be left for future work.

The AUG bolometry system¹² was used as a case study diagnostic system for testing the relative benefits of the volume ray-tracing technique although the conclusions should be generally applicable to other fusion devices.

II. RAY TRACING BOLOMETER MODEL

The total power (radiant flux) absorbed by a surface is given by the integral of the incident emission over the collecting solid angle Ω and surface area A ,

$$\Phi = \int_A \int_{\Omega} L_i(\mathbf{x}, \omega) \times \cos(\theta) d\omega dA. \quad (1)$$

Here, $L_i(\mathbf{x}, \omega)$ is the incident radiance arriving at a given point \mathbf{x} and an incident angle ω on the observing surface. The $\cos(\theta)$ term is a geometry factor describing the increase in the effective observing area as the incident rays become increasingly parallel to the surface.

Equation (1) is exact but extremely difficult to evaluate analytically for any realistic bolometer foil geometry and radiation distribution. In practice, it is easier to evaluate this integral with Monte Carlo integration and importance sampling which approximates the integral with a weighted average.^{13,14} The Monte Carlo integral estimator for a function f takes the form

$$I \approx \frac{1}{N} \sum_{j=1}^N \frac{f(x_j)}{p(x_j)}, \quad (2)$$

with $f(x)$ evaluated at N sample points x_j , and $p(x_j)$, given by

$$p(x_j) = \frac{q(x_j)}{\int q(x) dx}, \quad (3)$$

is the probability density function evaluated for the given sample point. $q(x)$ is the weight function for cases when the sample points are drawn from a non-uniform sample distribution.

The lighting integral in Eq. (1) can be naturally discretised in terms of N_r rays, composed of 2D sample points \mathbf{x}_j on detector area A_d and sample vectors ω_j on the hemisphere Ω . Therefore, the estimator for the power arriving on a bolometer foil would take the form

$$\Phi \approx \frac{1}{N_r} \sum_{j=1}^{N_r} \frac{L_i(\mathbf{x}_j, \omega_j) \cos(\theta_j)}{p_A(\mathbf{x}_j) p_{\Omega}(\omega_j)}. \quad (4)$$

If the sample points are drawn uniformly over the detector area, then $p_A(\mathbf{x}_j) = 1/A_d$. The natural choice for sampling the vectors is a uniform hemisphere. However, for a bolometer detector, the pinhole typically occupies a small solid angle leading to very computationally inefficient sampling. It is more efficient to sample the minimum cone of the solid angle (with half angle θ_h) that tightly wraps the pinhole. If the vectors are generated uniformly over the solid angle, the weighting function is still uniform, $q(\omega) = 1$, and thus, the probability density function takes the form of the fractional solid angle, $p_{\Omega}(\omega_j) = 1/(2\pi(1 - \cos(\theta_h))) = 1/\Omega_{frac}$. The estimator becomes

$$\Phi \approx \frac{\Omega_{frac} A_d}{N_r} \sum_{j=1}^{N_r} L_i(\mathbf{x}_j, \omega_j) \cos(\theta_j). \quad (5)$$

III. BENCHMARKING

To benchmark the implementation of Eq. (5), Raysect was configured to evaluate a set of test problems with known analytic answers. All cases use a simple spherical volumetric source as the emitter but demonstrate that the technique scales well to arbitrarily complex geometries.

Consider a sphere with a radius of 50 cm centred at the origin. If the sphere has a uniform radiance of $L_{sphere} = 1$ W/(m³/sr), then the total power radiated by the sphere's volume, V_{sphere} , is given by

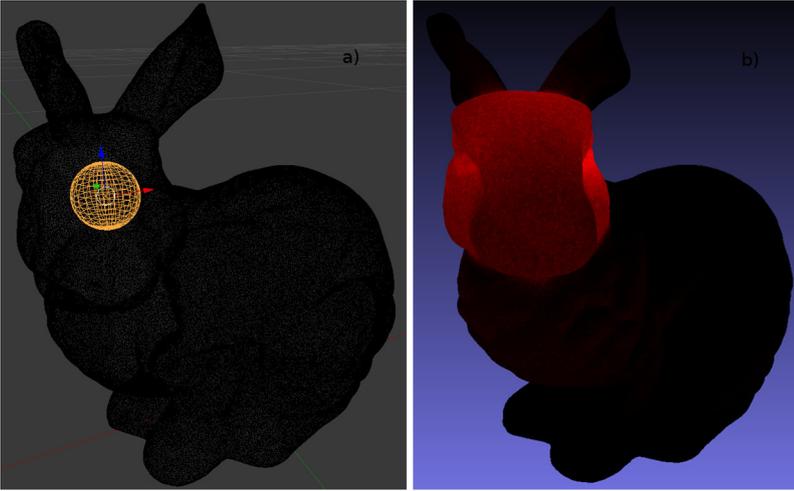


FIG. 2. (a) The Stanford Bunny mesh with a spherical light source of 1 cm radius positioned in the head region. (b) A log scaled power map of the power collected on the mesh surface elements.

$$\Phi_{theory} = L_{sphere} \times V_{sphere} \times 4\pi \approx 6.580 \text{ W}. \quad (6)$$

The radiating sphere can now be surrounded by an arbitrarily defined absorbing closed surface. As long as the surface does not intersect with the sphere, the total power collected by that surface through ray-tracing must equal the total power radiated, $\Phi_{ray-tracing} = \Phi_{theory}$. Changing the complexity of the confining surface provides a good benchmark since no matter how complex that surface becomes, the answer must always be the same.

Perhaps the simplest confining surface we can use is a cube centred on the origin with sides larger than the sphere's diameter. Note that due to symmetry, we do not need to observe each face of the cube, only a single face multiplied by 6. Test case A consisted of a cube with a side length of 2 m.

Case B considers a circular detector surface of radius 5 mm at a distance of 2 m from the spherical source used in case A. The expected power measurement is given by the total power radiated by the source multiplied by the fractional solid angle of the detector as seen by the source. This geometry approximates the scenario of a detector looking at a pinhole in the limit of large detector-pinhole separation.

Case C used one of the most commonly used test models in computer graphics, the Stanford Bunny.¹⁵ The Stanford Bunny is a mesh that was assembled from a set of range scans of a clay bunny. The mesh contains 69 451 triangles, forming a closed triangular mesh surface. The power load on each triangle in the mesh was individually ray-traced with Eq. (5), with the sum

over the mesh triangles giving the total power absorbed by the surrounding surface. For this test, the volumetric source was 1 cm in radius and approximately centred in the bunny's head region; see Fig. 2 for the source geometry.

As can be seen from Table I, in all benchmark cases, the Raysect ray-tracer recovers the analytical value to within the sampling uncertainty of the run. All computations were performed on a 16 core Intel Xeon E5-2665 at 2.4 GHz.

IV. ÉTENDUE CALCULATION

The étendue of the detector including occlusion and vignetting effects can be calculated by the weighted fraction of rays that pass through the slit multiplied by the full étendue that was sampled,

$$\epsilon_{det} = \frac{\Omega_{frac} A_d}{N_r} \sum_{j=1}^{N_r} \delta_j \cos(\theta_j), \quad \delta_j = \begin{cases} 0, & \text{ray hits} \\ 1, & \text{ray passes} \end{cases}. \quad (7)$$

Here, ϵ_{det} has units of [m² sr] and $\delta_j = 0$ if the j-th ray hits any of the obstructing aperture surfaces or $\delta_j = 1$ if the ray passes through unencumbered to the plasma.

With this calculation, the étendue for an arbitrary pinhole-foil geometry can be calculated to arbitrary precision with sufficient ray samples. By contrast, it is more common in bolometry to use the approximate pinhole formula^{2,4}

$$\epsilon_{pin} = \frac{\cos(\gamma) \cos(\alpha) A_d A_p}{d^2}, \quad (8)$$

where A_p is the rectangular pinhole area, A_d is the detector area, d is the distance between the pinhole and detector, and γ and α are compound angles describing the orientation of the foil surface with respect to the pinhole axis.^{2,4}

The limiting behavior of these two equations was explored with a simple pinhole geometry configuration. Consider a simple system where the detector and slit are both circular with a radius of $r = 1$ mm. The detector and slit are on axis and separated by some variable distance of separation, d , that scans from 0.1 mm up to 10 cm. The plot in Fig. 3 compares the ray-tracing results of Eq. (7) against the analytic approximation, Eq. (8).

TABLE I. Benchmarking of the implementation of ray-tracing Eq. (5) with case A: a cube with a side length of 2 m; case B: a pinhole camera-like geometry with large detector-pinhole separation; and case C: the Stanford Bunny mesh surface.¹⁵ The ray count, N_r , and computation time, t , are indicated for each case, along with the performance of the ray-tracing computation, $\Phi_{ray-tracing}$, against the analytic result, Φ_{theory} .

Case	N_r	Φ_{theory} (W)	$\Phi_{ray-tracing}$ (W)	t (s)
A	1×10^6	6.580	6.551 ± 0.057	0.05
B	1×10^6	1.028×10^{-5}	$1.030 \pm 0.003 \times 10^{-5}$	0.08
C	1×10^7	5.264×10^{-5}	$5.282 \pm 0.037 \times 10^{-5}$	1.13

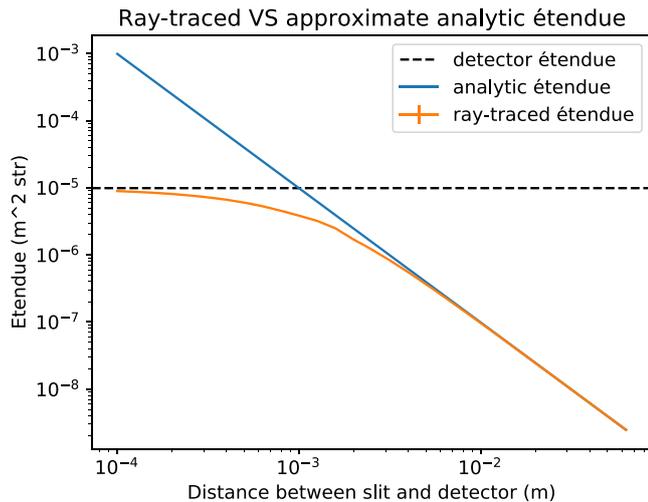


FIG. 3. The étendue of a simple pinhole camera composed of a circular detector and an aperture with radius 1 mm as a function of detector-aperture separation. The ray-traced and analytical étendues are plotted against each other, along with the limiting detector étendue.

As the separation tends to infinity, the ray-tracing and analytic formula agree to within the numerical sampling noise. In the opposite limit, the slit and detector eventually touch when the separation between them goes to zero. Since the detector and slit have the same geometrical shape and area, it is as if the slit was not present at all. And hence the étendue should tend toward the étendue of the detector on its own, represented by the dashed line in Fig. 3. The ray-tracing result correctly tends to this limit, while the analytic equation grows infinitely toward $d = 0$.

Figure 4 shows the relative error between the two calculation methods assuming that the ray-tracing method is more accurate. The figure shows the impact of changing the axis aligned aperture/detector geometry from circular to square and then increasingly elongated rectangles. The last data set includes an axis aligned aperture/detector combination with the geometry similar to the installed FLH bolometer camera

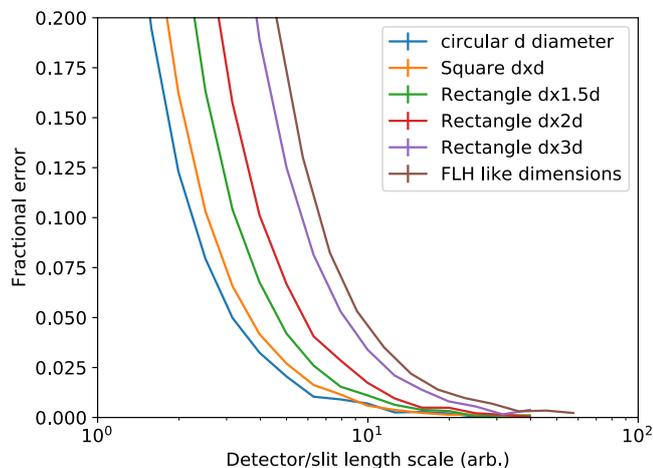


FIG. 4. The fractional error of the approximate analytical étendue relative to the ray-traced étendue for a number of different basic axis aligned aperture/detector geometries. The length scale is normalised to the smaller of the rectangular dimensions.

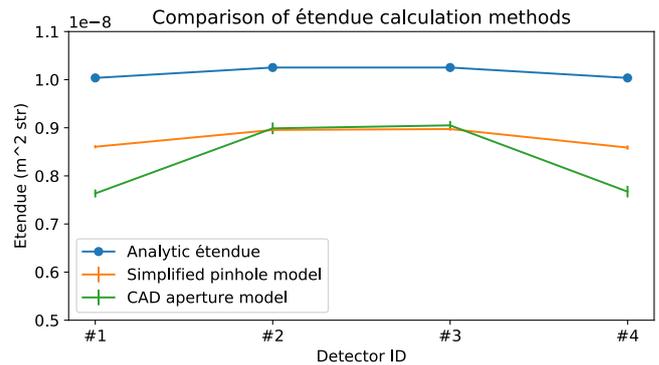


FIG. 5. The four foil detector étendues for the FLH camera were calculated with the analytic pinhole approximation [Eq. (8)], and these are compared with the ray-tracing étendue values [Eq. (7)]. The three cases shown are (a) the analytic pinhole approximation, (b) the ray-tracing calculation with a realistic first wall model and a simplified pinhole in a rectangular plane, and (c) as (b) but also including the full as-built detector geometry.

on AUG. Note that it is difficult to plot these curves on a normalised axis because of the subtle geometry differences between the different shapes. However, it demonstrates that the circular aperture is the limiting case and elongating the geometry leads to more deviation. Detectors and apertures are often elongated in the toroidal direction in fusion tomography systems to improve the signal to noise without significantly affecting the feature resolution. Although not shown, moving to a non-axis aligned system will further increase the amount of deviation.

The difference between the two étendue calculation methods was calculated for all AUG bolometer cameras. An extreme example case is shown for the FLH camera in Fig. 5. The ray-tracing calculation was performed with both a simplified [Fig. 5(b)] and full detail [Fig. 5(c)] aperture to separate out the different contributions. The simplified aperture used a rectangular plane with a rectangular pinhole cut out of the plane, whereas the full detail case used the engineering model of the as-built detector.

There is a systematic offset between the analytic (a) and ray-traced simple pinhole (b) cases of order $\sim 10\%$, which is due to performing the calculation with a simplified geometry. The extra drop in étendue between (b) and (c) is due to vignetting effects at the edges of the camera field of view when the full Computer Aided Design (CAD) geometry is included.

The detector étendue calculations for the FHC camera had an average offset of 10.8% compared with the much smaller 2.7% on the FLX camera. It was found that each camera had a distinctly different mean offset, which is likely due to differences in the camera geometry that have more or less effect on the assumptions in Eq. (8). The presence of an étendue offset will act as an extra systematic noise source in the inversion process.

V. SENSITIVITY MATRICES

Recovering the plasma emission with tomography is an ill-posed problem. It is customary to describe the system in terms of a sensitivity matrix \mathbf{W} . The elements $W_{k,l}$ describe the coupling between the N_s plasma emission sources x_l and

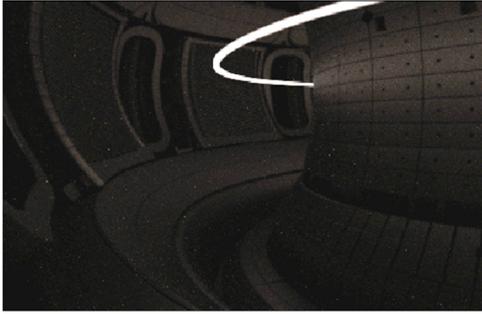


FIG. 6. A 3D voxel basis function composed of a toroidally symmetric annulus with a uniform volume emissivity. The full engineering model for the first wall is included in all ray-tracing calculations. Reflection effects have been turned on for visualisation purposes.

measured power Φ_k at N_d detectors. The whole detector set is typically represented as the matrix equation

$$\Phi = \mathbf{W}\mathbf{x}. \tag{9}$$

The power for the k th detector can be expressed as

$$\Phi_k = \sum_{l=1}^{N_s} W_{k,l} x_l, \tag{10}$$

where k and l are the indices for the detectors and sources, respectively. There are a number of possible choices for the prescription of the emitting source l basis functions.¹⁶ In this work, we have used a 3d voxel composed of a toroidally symmetric annulus with a uniform volume emissivity x_l ; see Fig. 6 for a visualisation of a single voxel element. Equation (9) can then be inverted using established tomography techniques^{2,3} to yield the spatial emissivities from a measured set of power values.

The conventional analysis technique estimates the single ray sensitivity matrix \mathbf{W}_{SR} by tracing a straight line through

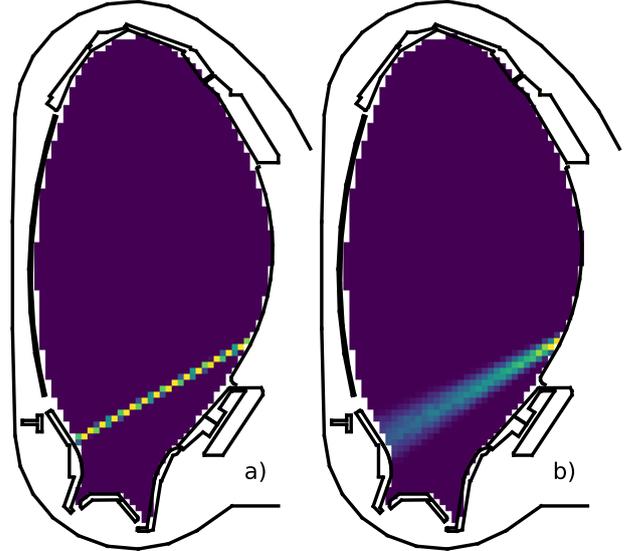


FIG. 7. Comparison of sensitivity matrices \mathbf{W} in the poloidal plane for a bolometer foil modelled with a single-ray and a volume sampled light cone.

the system, originating at the centre of the detector surface and passing through the slit centre. Every sensitivity element $W_{k,l}$ is weighted by the length of the ray segment, $s_{k,l}$, that intersects with the l th emitting source approximated as a 2d cell,

$$\mathbf{W}_{\text{SR}}: W_{k,l} = s_{k,l} \times \epsilon_{pin}. \tag{11}$$

By contrast, the ray-tracing technique launches N_r randomly generated rays from the detector and performs a weighted sum of the ray lengths, $s_{j,k,l}$, that intercept with the 3d voxel using Eq. (5). Rays that collide with the slit geometry or miss the cell have a zero weighting,

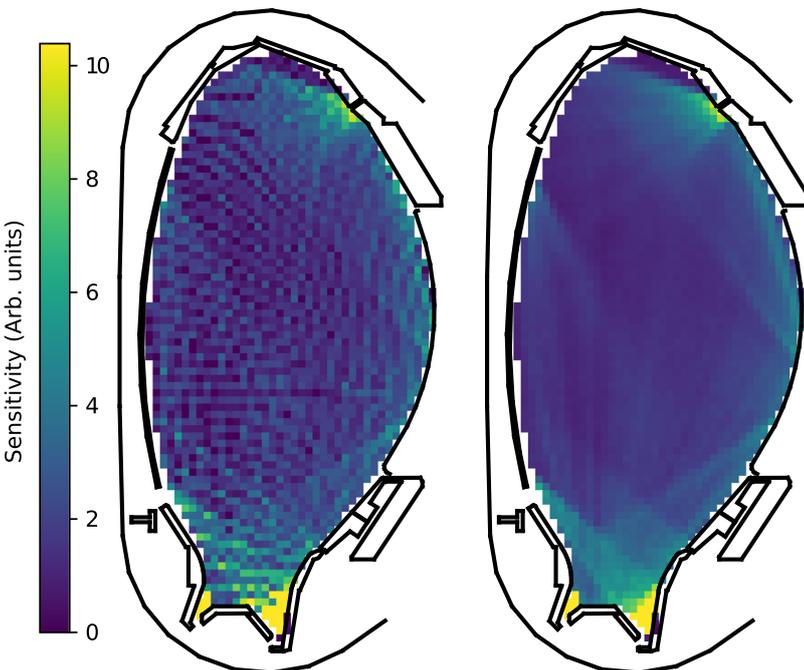


FIG. 8. Comparison of the sight line densities for foil bolometers at AUG modelled with single-ray paths and volume sampled light cones.

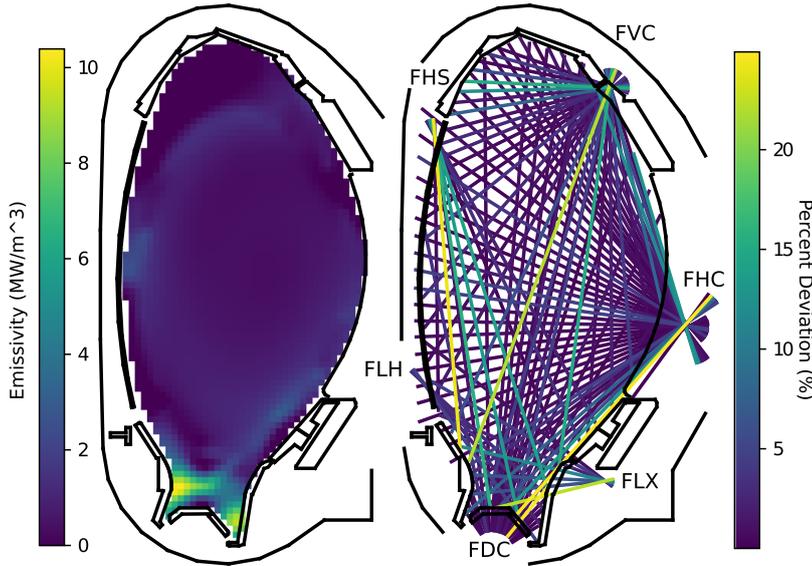


FIG. 9. (a) shows the BLB code inverted emission profile for AUG shot 33280 at 4.1 s. (b) shows the forward calculated sight lines colour-coded by the percentage error between the two techniques. The bolometer camera positions are labeled.

$$\mathbf{W}_{\text{Vol}} : W_{k,l} = \frac{\Omega_{\text{frac}} A_d}{N_r} \sum_{j=1}^{N_r} s_{j,k,l} \cos(\theta_j). \quad (12)$$

Figure 7 compares the sensitivity matrices for a single detector generated with the single ray and the volume ray-tracing methods. Unless the detector is perfectly collimated, the single ray method becomes less and less representative for voxels further from the detector. The coupling of a single ray to the voxels becomes more and more localised relative to the true extent of the detectors collection volume. By contrast, the volume sampling method leads to a smooth spatial response and the sensitivity region correctly expands as you get further from the detector. Additionally, the toroidal curvature of the voxels is automatically included due to the 3D nature of the calculation method.

A clear advantage of the volume sampling method is that the sensitivity matrices are less sensitive to design tolerances. Small changes in the position and direction of a single sight line can produce quite a big difference in the region of coupling for distant voxels. The volume ray-tracing matrix is much more stable to small perturbations in the input parameters.

Figure 8 compares the sight line density for the whole AUG foil bolometer detector set calculated with the two methods. The sight line density matrix is the normalised sum over the k index in \mathbf{W} , yielding a relative measure of how well a source region l is observed relative to the other cells. For the single ray technique, there are a number of cells in the plasma emission region that are effectively dark, i.e., not seen at all by a detector in the model. These cells can only be filled in by regularisation.

To study these effects further, we have used an example radiation scenario from AUG shot 33280 at 4.1 s inverted with the currently applied AUG tomography code,¹⁷ shown in Fig. 9(a). This tomography code uses the conventional single ray model for its inversions. Figure 9(b) shows each sight line colour coded by the percentage error in the observed power when the single-ray approximation is used in comparison with volume ray-tracing. Similarly, Fig. 10 shows a scatter plot of the forward modelled power calculated with a single-ray, Φ_{SR} ,

against the power calculated with a ray-traced volume, Φ_{Vol} , for each detector.

For many detectors observing lower powers, the single sight line is a good approximation; these sight lines are looking across the bulk plasma which is characterised by smoother emission gradients. However, the errors can become more significant ($>10\%$) for sight lines that see the divertor region with strong gradients, such as the strike point and x-point radiation. Single ray sight lines that tend to see a locally bright source will over-estimate the power, whereas sight lines that narrowly miss a bright source (such as a strike point) will under-estimate the power. It is also possible to under-estimate the power due to occlusion effects where the single ray path terminates too early compared to the full collection volume. An example of a sight line with this characteristic is shown in Fig. 11.

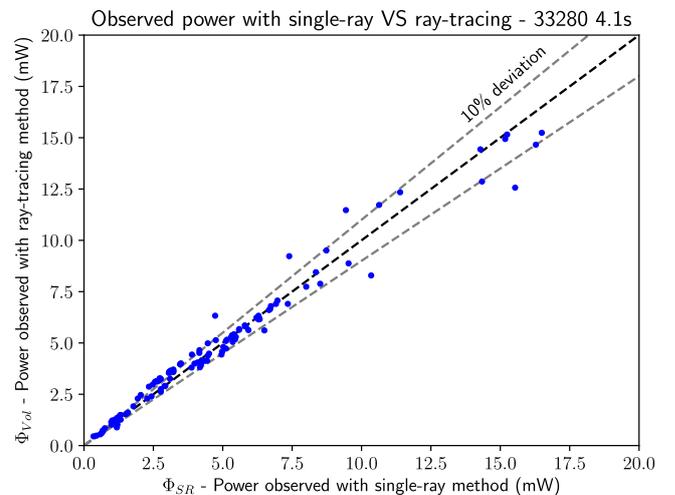


FIG. 10. Plot of the forward modelled power with the single-ray technique, Φ_{SR} , against the power calculated with ray-traced volumes, Φ_{Vol} , for each detector observing the radiation scenario in Fig. 9. The deviations become more pronounced at higher powers which tend to correlate with sight lines that see the divertor.



FIG. 11. Example case where the single-ray path terminates too early on a tile surface resulting in a significant error in the collected power calculation. When the volume ray-tracing technique is used, a large fraction of the collection volume extends into the inner divertor.

The sensitivity matrices for all AUG bolometer cameras were calculated with a 2206 cell inversion grid, taking between 1 and 4 h per camera on a 16 core Intel Xeon E5-2665 at 2.4 GHz.

VI. PERFORMANCE VS PHANTOMS

To explore the differences between reconstruction results obtained using the two sensitivity matrices, the two methods were tested on a standard population of 94 phantom emission scenarios¹⁸ used to benchmark the existing tomography code.¹⁷ The 94 phantoms are constructed from combinations of six basic emission sources: uniform backgrounds with gradients; point sources; x-points; strike points; divertor legs; and radiation rings on a flux surface.¹⁸ The full population of phantoms are designed to test the system's ability to resolve representative emission features that may be encountered in real plasmas.

Virtual observations for each phantom are constructed by multiplication of the volume sensitivity matrix with the phantom's emissivity vector, as per Eq. (9). In addition,

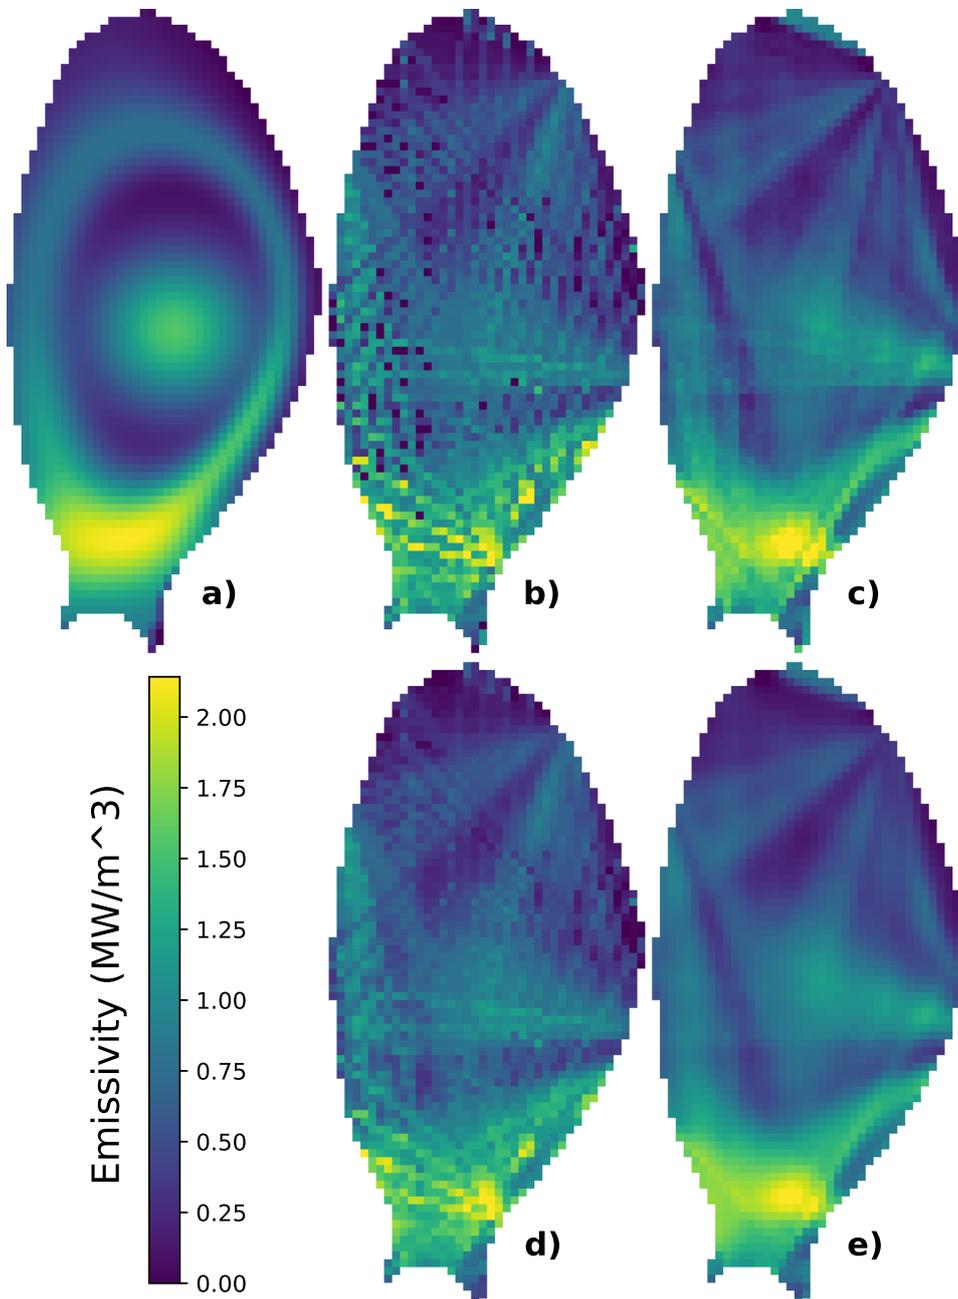


FIG. 12. Bolometer foil measurements of the phantom emission scenario given in (a) have been forward modelled with the volume ray-tracing method. The synthetic measurements are then inverted with the unregularised SART algorithm using sensitivity matrices constructed with the single-ray approximation (b) and the volume ray-tracing method (c). The differences between (b) and (c) demonstrate the extra spatial constraints imposed by volume ray-tracing. Additionally, both single-ray and ray-tracing inversions were used with regularised SART in (d) and (e), respectively. See Table II for a comparison of the results.

TABLE II. The total radiated power, P_{rad} , and Pearson correlation coefficient, ρ_c , for each of the inversions in Fig. 12. The regularised inversions have $\beta_L = 0.0125$.

	Phantom (a)	SR (b)	VOL (c)	SR + Reg (d)	VOL + Reg (e)
P_{rad} (MW)	13.82	13.37	13.76	13.45	13.7
ρ_c		0.68	0.85	0.79	0.89

5% Gaussian distributed noise was added to the virtual observations to simulate all uncertainties such as systematic errors in the detector alignment and calibration, as well as instrument noise in the detector electronics.

There are a wide range of inversion schemes used in fusion tomography diagnostics;² perhaps the most widely used is the Phillips-Tikhonov regularisation scheme combined with an anisotropic diffusion model as the objective function.¹⁷ In this work, the population of phantoms was inverted using the Simultaneous Algebraic Reconstruction Technique (SART)¹⁹ with an isotropic 2D discrete Laplacian smoothness operator. The strength of the smoothness hyper-parameter, β_L , could be scanned in order to quantify how far the results are independent of prior information. In addition, iterative schemes provide an easy way to enforce positivity since every iteration, the cells with negative emissivity can be clamped to zero. For more details on the SART implementation and regularisation parameters chosen, see the [Appendix](#).

To measure the performance of the inversions, we have opted for the Pearson correlation coefficient, ρ_c , which can measure the correlation between two vectors.²² In this context, it is defined as the covariance of the two emission source vectors, inversion \mathbf{x}_{inv} and phantom solution \mathbf{x}_{sol} , divided by the product of their standard deviations,

$$\rho_c = \frac{\text{Cov}(\mathbf{x}_{inv}, \mathbf{x}_{sol})}{\sigma(\mathbf{x}_{inv})\sigma(\mathbf{x}_{sol})}. \quad (13)$$

The SART inversions were performed on all 94 phantoms with sensitivity matrices computed with both the single-ray

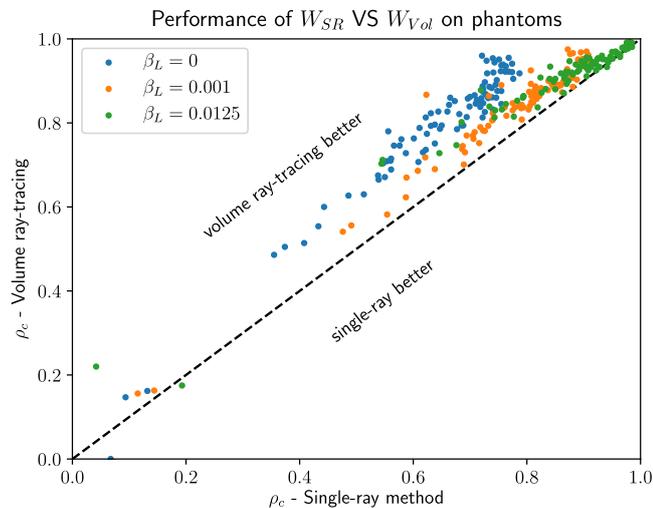


FIG. 13. Performance on all 94 phantoms on the two sensitivity matrices \mathbf{W}_{SR} and \mathbf{W}_{Vol} for varying levels of regularisation. The performance measure is the Pearson correlation coefficient, where $\rho_c = 1$ means that the inverted emission profile is identical to the phantom emission profile.

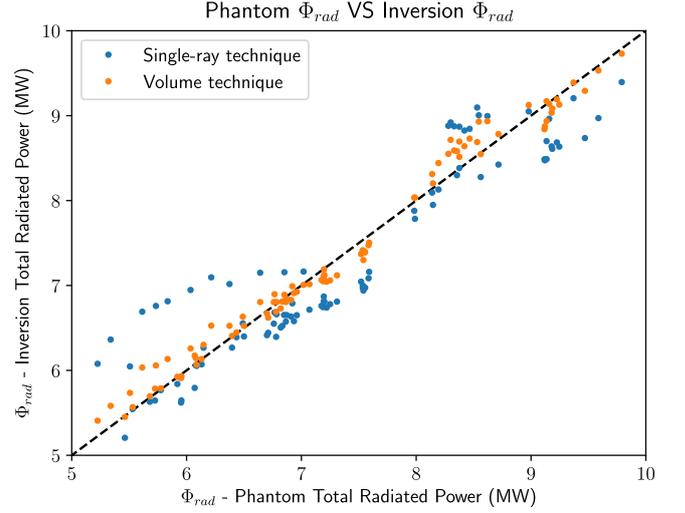


FIG. 14. The phantoms' total radiated power, Φ_{rad} , plotted against the inverted solution Φ_{rad} for all 94 phantoms using the single ray and volume ray-tracing techniques. Inversion performance decreases as the inversion points move away from the dashed line. This data set used the middle regularisation case ($\beta_L = 0.001$) and demonstrates that the volume ray-tracing technique consistently outperforms the single ray model.

(\mathbf{W}_{SR}) and volume ray-tracing (\mathbf{W}_{Vol}) techniques, both with and without the Laplacian gradient regularisation. The resulting inversions for an example phantom are shown in Fig. 12, with an accompanying comparison of the total radiated powers and correlation coefficients to the phantom in Table II. The differences between Figs. 12(b) and 12(c) demonstrate the extra spatial constraints imposed by the volume ray-tracing method. Figures 12(d) and 12(e) show the \mathbf{W}_{Vol} matrix continues to outperform \mathbf{W}_{SR} for the same level of gradient regularisation.

Figure 13 shows the performance on all phantoms for varying levels of regularisation. It demonstrates that the volume ray-tracing matrices are naturally more spatially constrained and require lower levels of regularisation for the same performance with single-rays. Furthermore, it is clear from Fig. 13 that turning up the regularisation is mainly affecting the single ray inversions. This is because the volume ray-tracing technique introduces fewer systematic errors due to being a more physically accurate model. At higher levels of regularisation, the two distributions converge as the regularisation starts to dominate the solution. This occurs because regularisation is a property of the solution matrix not the sensitivity matrix.

Figure 14 shows the distribution of the total radiated power for each inversion plotted against the phantom's real total radiated power. The distributions are plotted for both the single-ray and volume ray-tracing techniques using the middle regularisation case. The spread of points for the \mathbf{W}_{SR} case is bigger than the \mathbf{W}_{Vol} distribution with mean deviations of $5.2\% \pm 4.1\%$ and $1.7\% \pm 1.5\%$, respectively.

VII. DISCUSSION

The volume ray-tracing technique is expected to have the most impact on fusion machines where the solid angle of

the detectors is sufficiently large that the beam width of the detector sensitivity spans multiple plasma features or gradients. Therefore, this technique could have an impact on highly radiating divertor scenarios where the bulk of the radiation is local to the divertor and significant radiation gradients may be present.¹ The impact will be much less when the detector's sensitivity is highly collimated compared to the inversion grid.

The volume ray-tracing approach tends to produce sensitivity matrices where every grid cell is seen by multiple detectors. This has a natural smoothing effect on the solution matrix. The AUG bolometry system has a very high degree of spatial coverage and only ~5% dark cells in the inversion grid. But these effects can be more significant on other machines with poorer spatial coverage. Studies of the JET bolometry system came to similar conclusions and found that it was necessary to include the full geometry effects rather than relying on the line-integral approximation.⁴

Regularisation normally smoothens over many sources of error in an inversion process. Using volume ray-tracing to calculate the étendue and sensitivity matrices allows us to use a more physical model of the detector response and hence helps remove some of the systematic error sources. Therefore, it is not surprising that the volume ray-tracing approach has better performance at a given level of regularisation.

However, there are still other errors, such as detector noise and design errors that are always present in the data. There are also physical effects such as energy deposition/reflection of neutral particles and the finite foil reflectivity as a function of wavelength. The improvement due to volume ray-tracing degrades with increasing levels of noise where higher levels of regularisation are required. So it is important to have a system with a high level of signal to noise to see the benefits from this technique. On the other hand, there does not appear to be any disadvantages from including the extra calculation detail in the inversions.

The initial sensitivity matrix calculations for volume ray-tracing are more computationally intensive than for a single ray; however, they only need to be calculated once. Recalculation is only required when changing the inversion grid or the detectors, which is typically between experimental campaigns.

The volume ray-tracing method can be used with combined pinhole camera and first wall geometry to accurately model complicated compound apertures. This technique opens up a new design space for bolometer configurations that might not normally be considered, potentially helping to expand the tomography coverage in machines with poor accessibility.

VIII. CONCLUSIONS

A new technique for calculating the sensitivity of the bolometry detectors that accounts for the full geometric complexity has been presented. This technique was implemented using the Raysect open source ray-tracer and the CHERAB spectroscopy framework.

Volume ray-tracing techniques combined with full 3D machine models have been used to calculate the étendue

and sensitivity matrices of the bolometry system at ASDEX-Upgrade. The volume ray-traced sensitivity matrices were benchmarked against sensitivity matrices calculated with the more conventional single-ray technique. The volume ray-tracing technique can include vignetting and occlusion effects from installation features that are not possible to include in simpler calculation methods.

Inversions using both matrices were carried out on a population of emission phantoms with varying levels of regularisation. The volume ray-tracing technique consistently outperformed the single-ray technique in regard to both the correlation coefficient distance measure and the accuracy of the inverted total radiation power. The volume ray-tracing technique was shown to be naturally more spatially constrained.

The technique is a useful addition to the standard bolometry techniques currently in use and does not present any obvious disadvantages. However, the volume ray-tracing method requires detailed *in situ* information of the system and the improvements gained from including the extra geometric detail degrade quickly with increasing sources of error such as detector noise and design errors. It is important for a bolometry system to have good signal to noise ($\leq 5\%$ detector noise) to gain significant benefit from the extra calculation detail.

ACKNOWLEDGMENTS

The authors wish to acknowledge the contributions of V. Brack at IPP Garching who extracted the relevant bolometer camera geometry from the engineering CAD files.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014–2018 under Grant Agreement No. 633053 and from the RCUK Energy Programme (Grant No. EP/P012450/1). The views and opinions expressed herein do not necessarily reflect those of the European Commission.

B. Lipschultz was funded in part by the Wolfson Foundation and UK Royal Society through a Royal Society Wolfson Research Merit Award as well as by the RCUK Energy Programme (EPSRC Grant No. EP/I501045).

APPENDIX: TOMOGRAPHIC INVERSION METHOD

The population of phantoms were inverted with the Simultaneous Algebraic Reconstruction Technique (SART).¹⁹ The SART method is an iterative inversion scheme where the emission cells are updated with the formula

$$x_l^{(i+1)} = f_{sart}(x_l^{(i)}) = x_l^{(i)} + \frac{\omega}{W_{\oplus,l}} \sum_{k=1}^{N_d} \frac{W_{k,l}}{W_{k,\oplus}} (\Phi_k - \hat{\Phi}_k), \quad (\text{A1})$$

where

$$W_{k,\oplus} = \sum_{l=1}^{N_s} W_{k,l}, \quad W_{\oplus,l} = \sum_{k=1}^{N_d} W_{k,l}.$$

Here $x_l^{(i)}$ is the previous estimate for the emission at source l in iteration i . The relaxation hyperparameter was set to $\omega = 1$

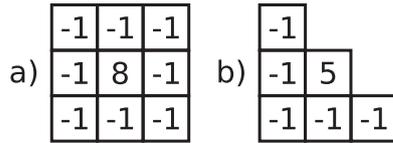


FIG. 15. Examples of the 2D Laplacian operator for (a) a central cell ($C = 8$) and (b) a cell near the inversion grid corner ($C = 5$).

for this study. The SART method effectively updates each cell by the weighted average error between the forward modelled $\hat{\Phi}_k$ and observed Φ_k measurements. The observed errors are weighted by both their proportion of the total ray length ($W_{k,\oplus}$) and the sum of the effective ray paths crossing that cell ($W_{\oplus,l}$).

Because the inversion process is an ill-posed problem, there are technically an infinite set of solutions to Eq. (9). The sensitivity matrix \mathbf{W} can have a non-trivial null space, meaning that any linear combination of null space models can be added to a particular solution and not change the fit to the data. Regularisation is a technique that applies additional constraints in the form of an objective function to bias the inversion process toward solutions that satisfy our prior knowledge.^{2,21} The objective function is a regularisation operator that quantifies some properties of the emission profile such as Laplacian smoothness, anisotropic smoothness, or minimum cross-entropy.

The current AUG code employs the Anisotropic Diffusion Model Tomography¹⁷ (ADMT) scheme, which assumes that radiation varies less along magnetic flux surfaces than perpendicular to them. This scheme uses an anisotropic smoothness operator and is prescribed in terms of two diffusion terms, perpendicular D_{\perp} and parallel D_{\parallel} to flux surfaces. In this work, we have instead opted for a isotropic 2D discrete Laplacian smoothness operator,²⁰

$$\hat{\mathcal{L}}_{iso}(x_l^{(i)}) = \beta_L (Cx_l^{(i)} - \sum_{c=1}^C x_c^{(i)}). \quad (\text{A2})$$

Here, c is the index for the sum over the eight possible neighbouring cells. This regularisation operator was chosen to reduce the amount of prior information. See Fig. 15 for an example of the operator in two different grid positions. With the Laplacian smoothness objective, the update formula becomes

$$x_l^{(i+1)} = f_{sart}(x_l^{(i)}) - \hat{\mathcal{L}}_{iso}(x_l^{(i)}). \quad (\text{A3})$$

There are two regularisation hyperparameters in this scheme. The first is β_L which determines the amount of local smoothness imposed. When $\beta_L = 0$, the solution is fully determined by the measurements, and as $\beta_L \rightarrow 1$, the solution is dominated by the smoothness operator. The criterion for choosing the value of β_L can be determined by a constraint such as the misfit between the forward modelled $\hat{\Phi}_k$ and observed Φ_k measurements or the L-curve method.^{21,23} In this work, the β_L parameter was scanned in order to quantify how far the results are independent of prior information.

The other regularisation parameter is the criterion for terminating the SART iterations. In the implemented scheme, the iteration loop was broken when either the difference between successive χ^2 values fell below 10^{-4} or a maximum of 250 iterations was reached.

After every iteration, all cells with negative emissivity were clamped to zero to enforce positivity.

In future work, the isotropic smoothness operator could be exchanged for two 2D discrete gradient operators aligned parallel and perpendicular to the local magnetic field, $\hat{\mathcal{L}}_{\parallel}$ and $\hat{\mathcal{L}}_{\perp}$, with accompanying hyperparameters β_{\parallel} and β_{\perp} . These anisotropic operators would have an effect similar to the established ADMT scheme¹⁷ and allow further assessment of the value of including the extra geometric detail when known physics is present.

¹F. Reimold, M. Wischmeier, M. Bernert, S. Potzel, A. Kallenbach, H. W. Müller, B. Sieglin, U. Stroth, and ASDEX Upgrade Team, *Nucl. Fusion* **55**, 033004 (2015).

²L. C. Ingesson, B. Alper, B. J. Peterson, and J. C. Vallet, *Fusion Sci. Technol.* **53**, 528 (2008).

³T. Craciunescu, G. Bonheure, V. Kiptily, A. Murari, I. Tisceanu, V. Zoita, and JET-EFDA Contributors, *Nucl. Instrum. Methods Phys. Res., Sect. A* **605**, 374 (2009).

⁴L. C. Ingesson, C. F. Maggi, and R. Reichle, *Rev. Sci. Instrum.* **71**, 1370 (2000).

⁵L. C. Ingesson, P. J. Böcker, R. Reichle, M. Romanelli, and P. Smeulders, *J. Opt. Soc. Am. A* **16**, 17 (1999).

⁶L. C. Ingesson and D. J. Wilson, *Rev. Sci. Instrum.* **73**, 2890 (2002).

⁷M. Weiland *et al.*, *Plasma Phys. Controlled Fusion* **57**, 085002 (2015).

⁸D. Vezinet, V. Igochine, M. Weiland, Q. Yu, A. Gude, D. Meshcheriakov, M. Sertoli, ASDEX Upgrade Team, and EUROfusion MST1 Team, *Nucl. Fusion* **56**, 086001 (2016).

⁹C. Giroud, A. Meakins, M. Carr, A. Baciero, and C. Bertrand (2018). "CHERAB spectroscopy modelling framework," Zenodo, V0.1.0, Dataset <http://doi.org/10.5281/zenodo.1206142>.

¹⁰M. Carr *et al.*, "Towards integrated data analysis of divertor diagnostics with ray-tracing," in 44th EPS Conference on Plasma Physics, Belfast, Northern Ireland, UK, 26–30 June 2017 (<http://ocs.ciemat.es/EPS2017PAP>).

¹¹A. Meakins and M. Carr (2017). "Raysect Python raytracing package," Zenodo, V0.4.0, Dataset <http://doi.org/10.5281/zenodo.1205064>.

¹²M. Bernert, T. Eich, A. Burckhart, J. C. Fuchs, L. Giannone, A. Kallenbach, R. M. McDermott, B. Sieglin, and ASDEX Upgrade Team, *Rev. Sci. Instrum.* **85**, 033503 (2014).

¹³M. Pharr and G. Humphreys, *Physically Based Rendering: From Theory to Implementation* (Morgan Kaufmann, 2016).

¹⁴E. Veach, "Robust Monte Carlo methods for lighting simulation," Ph.D. thesis, Stanford University, 1997.

¹⁵G. Turk and M. Levoy, "Zipped polygon meshes from range images," in *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (ACM, 1994).

¹⁶L. C. Ingesson, H. Chen, P. Helander, and M. J. Mantinen, *Plasma Phys. Controlled Fusion* **42**, 161 (2000).

¹⁷J. C. Fuchs *et al.* in *Europhysics Conference Abstracts: 21st Conference on Controlled Fusion and Plasma Physics, Montpellier, June 27* (European Physical Society, 1994), Vol. 18B, p. 1308.

¹⁸S. Glöggler, "Development of an algorithm to determine the total radiated power," Max-Planck-Institut für Plasmaphysik Report No. IPP 2018-06, 2018.

¹⁹A. Andersen and A. Kak, *Ultrason. Imaging* **6**, 81 (1984).

²⁰T. Hobiger, T. Kondo, and Y. Koyama, *Earth, Planets Space* **60**, 727 (2008).

²¹R. Aster, B. Borchers, and C. Thurber, *Parameter Estimation and Inverse Problems* (Academic Press, 2011), Vol. 90.

²²K. Pearson, *Proc. R. Soc. London* **58**, 240 (1895).

²³P. C. Hansen, *SIAM Rev.* **34**, 561 (1992).