

3DMPR – a robust morphological approach for applying phase retrieval in proximity to highly attenuating objects in computed tomography

James A. Pollock,^{a*} L. C. P. Croton,^a K. S. Morgan,^a K. J. Crossley,^{b,c}
M. J. Wallace,^{b,c} G. A. Buckley,^d S. B. Hooper,^{b,c} and M. J. Kitchen^{a*}

Received 24 October 2024

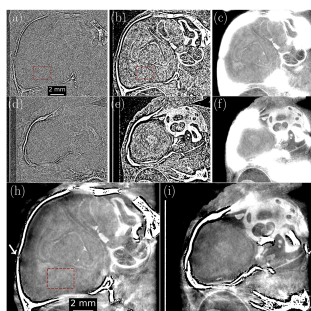
Accepted 27 June 2025

Edited by M. Guizar-Sicairos, Paul Scherrer Institute, Switzerland, and EPFL, Switzerland

Keywords: X-ray imaging; phase contrast; high contrast; multi-material samples; synchrotron imaging; computed tomography; phase retrieval; image processing; image segmentation.

^aSchool of Physics and Astronomy, Monash University, Clayton, VIC, Australia, ^bRitchie Centre, Hudson Institute of Medical Research, Clayton, VIC, Australia, ^cDepartment of Obstetrics and Gynaecology, Monash University, Wellington Rd, Clayton, VIC, Australia, and ^dWalter and Eliza Hall Institute of Medical Research, Parkville, VIC, Australia. *Correspondence e-mail: james.pollock@monash.edu, marcus.kitchen@monash.edu

X-ray imaging is a fast, precise and non-invasive method of imaging which, when combined with computed tomography, provides detailed 3D rendering of samples. Incorporating propagation-based phase contrast can vastly improve data quality for weakly attenuating samples via phase retrieval, allowing radiation exposure to be reduced. However, applying phase retrieval to multi-material samples commonly requires the choice of which material boundary to tune the reconstruction. Selecting the boundary with strongest phase contrast increases noise suppression, but at the detriment of over-blurring other interfaces and potentially removing quantitative sample information. Additionally, conventional phase retrieval algorithms cannot be used for regions bounded by more than one material, requiring alternative methods. Here we present a computationally efficient, non-iterative nor AI-mediated method for applying strong phase retrieval, whilst preserving sharp boundaries for all materials within the sample. 3D phase retrieval is combined with morphological operations to prevent over-blurring artefacts from being introduced, while avoiding the potentially long convergence times required by iterative approaches. This technique, entitled 3DMPR, was tested on phase contrast images of a rabbit kitten brain encased by the surrounding dense skull. Using 24 keV synchrotron radiation with a 5 m propagation distance, 3DMPR provided a 6.8-fold improvement in the signal-to-noise ratio (SNR) of brain tissue over the standard phase retrieval procedure, without over-smoothing the images. Simultaneous quantification of edge resolution and SNR gain was performed with an aluminium–water phantom imaged using a microfocus X-ray tube at 35 kV_p and 0.576 m effective propagation distance. There, 3DMPR provided a four-fold SNR boost whilst preserving the boundary spatial resolution at $54 \pm 1 \mu\text{m}$, compared with $108 \pm 2 \mu\text{m}$ using conventional phase retrieval. These results illustrate the ability of 3DMPR to create new avenues of dose reduction in clinical settings.



1. Introduction

X-ray imaging can provide high-resolution, three-dimensional (3D) visualization of the internal structure of an object when combined with computed tomography (CT). However, capturing high-resolution CT images with high signal-to-noise ratio (SNR) using standard X-ray systems requires a non-negligible radiation dose, potentially harming samples that are sensitive to ionizing radiation. Hospital imaging systems limit radiation exposure by using highly efficient X-ray detectors, with relatively low resolution being the typical trade-off. Iterative reconstruction algorithms may also be used to help balance noise and spatial resolution (Vliegthart *et al.*, 2022).

High-resolution CT scans of biological tissues can be achieved with low radiation dose by using phase contrast X-ray imaging (Kitchen *et al.*, 2017; Arhatari *et al.*, 2021). Phase contrast imaging uses specific setups to transform phase effects, introduced by the sample, into intensity variations capable of being recorded by detectors, and hence provides more information than from absorption contrast alone. While several X-ray phase contrast techniques exist, we focus here on propagation-based imaging (PBI), taking advantage of the simplicity in optical design that only requires a sufficiently coherent (Pelliccia *et al.*, 2018) X-ray wavefield at the sample position and some propagation distance between the sample and detector. PBI has relaxed temporal coherence requirements and hence is applicable to macro- (Kitchen *et al.*, 2017; Pollock *et al.*, 2025) and micro-scale (Bidola *et al.*, 2017) imaging with synchrotron and laboratory sources. Phase contrast arises from interference between sections of the X-ray wavefield that have incurred different phase shifts as they pass through an object. This results in intensity fringes appearing along material boundaries at the detector plane, effectively acting as a sharpening filter (Gureyev *et al.*, 2017). This is particularly valuable for those boundaries between low- Z materials, which are weakly attenuating and hence create only weak attenuation contrast. To restore the sample structure requires the application of a phase retrieval algorithm. For objects comprising a single monomorphous material, the algorithm of Paganin *et al.* (2002) accomplishes this by inverting the transport-of-intensity equation (TIE). The result is a low-pass filter specifically tuned to suppress phase contrast fringes at the boundaries of the object, whilst also providing SNR amplification by suppressing high-frequency noise. The algorithm can be represented by

$$I(x, y) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F} [I_d(x, y)/I_0]}{1 + \frac{\delta\Delta}{\mu} \mathbf{k}_\perp^2} \right\}, \quad (1)$$

where \mathcal{F} and \mathcal{F}^{-1} are the forward and inverse Fourier transform operators, $I_d(x, y)$ is the image measured at the detector plane, I_0 is the spatially variable incident intensity, Δ is the propagation distance from object to detector, δ represents the real difference of the complex refractive index from unity and μ is the linear attenuation coefficient. \mathbf{k}_\perp^2 describes the transverse spatial frequency components in order to provide variable weighting to the low-pass Lorentzian filter. While equation (1) applies to the two-dimensional projection images, phase retrieval can similarly be performed after CT reconstruction following

$$\mu_{\text{PR}}(x, y, z) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F} [\mu_{\text{PC}}(x, y, z)]}{1 + \frac{\delta\Delta}{\mu} \mathbf{k}_{3\text{D}}^2} \right\} \quad (2)$$

(Thompson *et al.*, 2019) where $\mathbf{k}_{3\text{D}}^2 = k_z^2 + k_y^2 + k_x^2$, $\mu_{\text{PC}}(x, y, z)$ is the CT volume of linear attenuation coefficients reconstructed from flat-fielded phase contrast projections, and μ_{PR} is the same CT volume after application of phase retrieval. Although applying phase retrieval in 3D allows slightly higher SNR gains compared with equation (1), each method is a fast,

stable and effective means of achieving low-dose discrimination between low- Z monomorphous materials and their boundaries with vacuum or air (Kitchen *et al.*, 2017; Beltran *et al.*, 2011). Any additional material interfaces present in the object will either be under-blurred, resulting in remnant phase contrast fringes or, in the case of higher- Z materials, over-blurred, distorting the boundary and leaving affected regions non-quantitative or completely obscured (Beltran *et al.*, 2010; Beltran *et al.*, 2011; Croton *et al.*, 2018a). To illustrate this, we used the projection approximation to simulate the exit surface wavefield of a low- Z object, embedded with spherical cavities and a high- Z cylinder, shown in Fig. 1(a). This wavefield was then propagated using the transport of intensity equation (TIE) to produce phase contrast fringes on all material boundaries [Fig. 1(b)]. Gaussian noise was added to demonstrate the benefit of phase retrieval, which still accurately restores the low- Z cavities [Fig. 1(c)], although at the expense of over-blurring the high- Z interface.

A later generalization of equation (1) allows the retrieval to be specifically tuned to a specific interface within the object between two different materials (Beltran *et al.*, 2010; Croton *et al.*, 2018a), introducing a second material component in δ_2 and μ_2 as

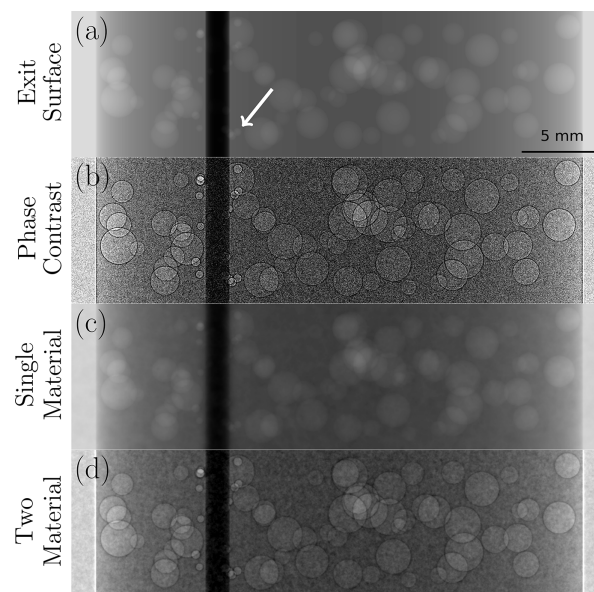


Figure 1 Simulations of phase retrieval algorithms applied to projections of a two-material phantom. This phantom is composed of low- Z PMMA embedded with spherical cavities (white arrow) and a high- Z aluminium cylinder which appears as a dark column in the exit surface wavefield shown in (a). Panel (b) shows the phase contrast image produced from free-space propagation via the TIE to produce Fresnel fringes at each boundary that vary in amplitude according to the materials present. Here, noise is added to simulate low-flux imaging. (c) Single material phase retrieval with equation (1), using δ and μ for PMMA, showing accurate restoration of the low- Z /air boundaries [*cf.* panel (a)] but severely over-blurs the high- Z material and the surrounding area. This can be avoided with (d) two-material phase retrieval to the high- Z , low- Z interface, using equation (3) with the δ and μ values for the low- Z (PMMA) and high- Z (Al) materials, but at the expense of higher remnant noise and phase contrast fringes left around the low- Z /air boundaries of each cavity. An ideal approach would achieve the benefits of both approaches shown in (c) and (d).

$$I(x, y) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F} [I_d(x, y)/I_0]}{1 + \frac{(\delta_2 - \delta_1)\Delta}{\mu_2 - \mu_1} \mathbf{k}_\perp^2} \right\}. \quad (3)$$

This allows phase retrieval to be applied to any interface between two adjacent, non-air materials, reconstructing that edge correctly. However this approach leaves any other interfaces either over-blurred or with remnant phase contrast fringes, as shown in Fig. 1(d), around the air cavities. If the material properties are similar [*i.e.* $(\delta_2 - \delta_1)/(\mu_2 - \mu_1)$ is small], the phase contrast is weak, so the phase retrieval filter is equivalently weak and only lightly suppresses image noise.

Beltran *et al.* (2011) showed that, if specific pairs of material interfaces exist in the object, a complete CT reconstruction can be comprised without over- or under-smoothed boundaries by performing separate phase retrieved CT reconstructions for each material pair, then interleaving those reconstructions together. That technique required large spatial tolerances around high- Z components to avoid including over-blurred regions in the final stitched image. Hence, it required features of interest to be well spatially separated from high- Z materials, so that strong phase retrieval filtering did not spread intensity into the neighbouring feature of interest – referred to as over-blurring artefacts. Alternatively, iterative approaches aim to avoid this problem by masking out the high- Z material from projection images and using feedback cycles from the CT reconstructions to correct interface over-blurring caused by phase retrieval (Hehn *et al.*, 2018a; Hehn *et al.*, 2018b). The present paper presents an alternative method that combines 3D phase retrieval with morphological operations to prevent over-blurring artefacts from being introduced and avoids the potentially long convergence times required of iterative approaches. Hereafter, this technique will be referred to as 3D Masked Phase Retrieval (3DMPR).

Typically, phase retrieval applications in pre-clinical studies focus on low- Z tissues where the SNR gains are highest, as seen in synchrotron studies of the breast (Gureyev *et al.*, 2014; Nesterets *et al.*, 2015) and lungs (Pollock *et al.*, 2025; Albers *et al.*, 2023; Kitchen *et al.*, 2017; Mohammadi *et al.*, 2014; Dullin *et al.*, 2015), or recent small animal imaging applications on freeze-dried hearts using micro-focus sources (Lioliou *et al.*, 2024). Their mutual aim is to improve resolution and SNR while reducing radiation dose imparted to potential patients. Similar benefits have been demonstrated in brain imaging using synchrotron radiation (Croton *et al.*, 2018a; Croton *et al.*, 2018b), although phase retrieval has been restricted through equation (3) to the bone–tissue interface as the surrounding strongly attenuating skull can easily be over-blurred by strong phase retrieval filtering, obscuring the periphery of the brain. X-ray CT reconstructions are pertinent to medical applications for their advantages over MRI scanners such as suitability for evaluating penetrating brain injuries and acute neurological emergencies in time-critical situations (Temple *et al.*, 2015; Brody *et al.*, 2015). Successful implementation of 3DMPR to brain imaging would open additional avenues for pursuing low-dose preclinical trials for *in situ* brain imaging and further secure PBI phase contrast imaging as a safe and reliable means

of high speed and high resolution brain imaging. Additionally, phase retrieval could be made compatible with standard contrast-agent data (Lang *et al.*, 2014; Lang *et al.*, 2016; Reichmann *et al.*, 2023), expanding the flexibility of PBI imaging. 3DMPR could also be applied to the Human Organ Atlas project for regions where soft tissue meets highly attenuating bone.

Section 2 describes the new 3D computational process for applying strong phase retrieval to objects that contain both weakly and highly absorbing materials. Section 3 shows examples of multi-material CT reconstructions performed with 3DMPR, including for polychromatic cone-beam data using a microfocus X-ray source, while providing resolution and SNR comparisons to where conventional phase retrieval is applied in 2D prior to reconstruction.

2. Methods

This section provides a complete description of the computational process, demonstrated using the propagation-based phase contrast CT dataset of the head of a rabbit kitten recorded at beamline 20B2 of the SPring-8 synchrotron, Japan. Henceforth, to keep the method description general, the low-effective- Z brain-tissue, composed of atomic elements H:C:N:O:Na:P:S:Cl:K in the approximate ratio 8510:968:126:3567:7:10:5:7:6 with a density of 0.986 g cm^{-3} (Chantler *et al.*, 2005; Berger *et al.*, 2010), will be referred to as material A. Similarly, the high-effective- Z bone surrounding the brain, H:C:N:O:Na:Mg:P:S:Ca in the ratio 3878:1483:345:3125:5:9:11:11:645 with density 1.45 g cm^{-3} , will be referred to as material B. When applying this method to other samples, the more attenuating material within the sample will be designated as material B.

Imaging was performed using synchrotron radiation at 24 keV, with a 5 m propagation distance, and recorded using a 2048×2048 pixel Hamamatsu sCMOS camera (C11440-22C) with $6.5 \mu\text{m}$ pixel size, fibre-optically coupled to a $15 \mu\text{m}$ -thick Gadox ($\text{Gd}_2\text{O}_2\text{S:Tb}^+; \text{P43}$) phosphor. Due to the small width of this detector, the object was positioned with one half filling the detector and rotated through 360° so that all object features are moved through the field of view. Antipodal images of the CT dataset are then stitched together, resulting in a new dataset of the whole object spinning through 180° and separated by 0.1° increments, now with double the width. These CT datasets were reconstructed with a parallel beam geometry using filtered back-projection through the *TomoPy* and *Astra* packages (van Aarle *et al.*, 2016; van Aarle *et al.*, 2015) in Python, which was the language used for all analysis steps. Material parameters of the complex refractive index were taken as $\mu = 55.1 \text{ m}^{-1}$, $\delta = 3.93 \times 10^{-7}$ for material A, and $\mu = 336.83 \text{ m}^{-1}$, $\delta = 5.43 \times 10^{-7}$ for material B (Schoonjans *et al.*, 2011).

Fig. 2 provides a flow diagram of all the steps required to produce the phase retrieved output, beginning with the phase-contrast volume represented by the example CT slice in Fig. 2(a). Although all stages of the method are represented via the same CT slice, we note that 3D phase retrieval is a

volume operation requiring sample volumes of dimensions at least as large as the blurring kernel. In addition to the raw phase-contrast CT volume (*i.e.* reconstructed without phase retrieval), we require a CT volume that has been phase retrieved for the A/B material interface, using equation (3). This volume, represented by Fig. 2(b), can be calculated by phase retrieving the phase-contrast volume in 3D (Thompson *et al.*, 2019) using equation (2), or by creating a new CT reconstruction from phase retrieved projections (Beltran *et al.*, 2010).

The phase retrieved volume is used to produce a binary array of the material B locations using a simple threshold that is expanded with a 3D dilation filter to ensure all voxels containing material B are included. Fig. 2(c) shows an example binary array, calculated using a threshold of $\mu = 77.5 \text{ m}^{-1}$ and requiring 16 iterations of a $3 \times 3 \times 3$ kernel dilation filter due to the presence of the low absorption region (dark region) directly inside the skull. This kernel was chosen since it symmetrically expands the mask with a one pixel border; however, this could be varied as required to ensure complete coverage of all material B pixels. Using the binary array, Fig. 2(c), allows material B to be masked out of the phase-contrast volume, Fig. 2(a), replacing the grey values of those voxels with the theoretical attenuation coefficient of material A. This results in Fig. 2(d). Previously, we attempted to fill the masked pixels using interpolation algorithms; however, even the levels of noise in the images retrieved using equation (3) created too much instability, causing high contrast boundaries to appear in homogeneous sample regions. Substitution of a constant value instead minimizes local contrast in and near the masked region of the image and only requires prior knowledge of the low-Z material, an assumption already required by

the phase retrieval algorithms. Alternatively, the replacement attenuation coefficient could be manually measured from the phase retrieved reconstruction [Fig. 2(b)]. Next, we apply 3D phase retrieval to the masked phase contrast volume, following equation (2), to produce the low-noise phase retrieved CT volume in Fig. 2(e). Having removed any high-Z material from our sample, this retrieval cannot create any over-blurring artefacts, and performing the phase retrieval in 3D additionally provides a slight SNR boost over projection-based phase retrieval (Thompson *et al.*, 2019).

Although 3D phase retrieval of large volume arrays can be RAM intensive, we were able to implement an out-of-core approach on a Dell Precision 5560 laptop with 32 GB RAM, taking advantage of high SSD read and write speeds. Phase retrieval of a $1030 \times 1030 \times 1030$ voxel data volume was achieved in just 4.6 min. Advanced implementation on high RAM devices/computer clusters will reduce calculation times but is generally not required. After the mask threshold and dilation settings have been selected for a particular material combination, reconstructions of similar composition could be calculated consecutively without further manual intervention, providing the possibility of rapid feedback during an experiment.

Finally, we use the material B binary array [Fig. 2(c)] to interleave the altered regions of the phase retrieved volume, Fig. 2(e), with the A/B phase retrieved regions of Fig. 2(b). This completes our method of 3DMPR, producing Fig. 2(f), a low-noise CT volume that benefits from the strong phase retrieval designed for a low contrast material to within a few pixels of the high-Z material interface, without introducing over-blurring artefacts. This provides noise reduction within material A without eliminating features close to the A/B

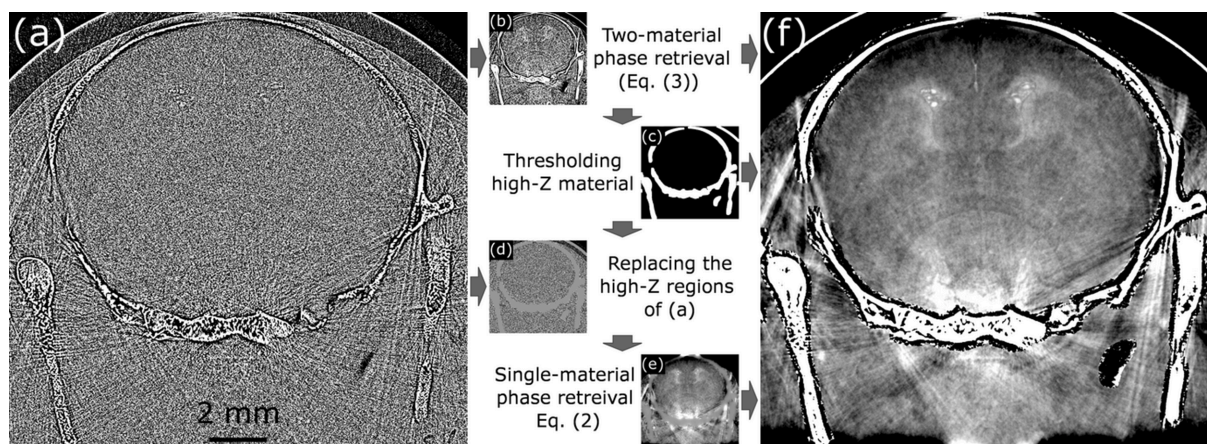


Figure 2 The masking and phase retrieval process, using a rabbit brain cross-section for demonstration. Images are labelled alphabetically according to the order in which they are computed while arrows indicate precursor images required to calculate the subsequent image. (a) A phase-contrast CT slice without any reprocessing applied, beyond flat and dark correction. (b) The same slice phase retrieved for the bone/brain tissue interface [equation (3)], either achieved from 3D phase retrieval of (a) or CT reconstruction of phase retrieved projections. Upon thresholding (b), we create a pixel-wise binary mask of the high-Z material, applying a dilation filter to ensure full enclosure of the high contrast boundaries. The binary mask, (c), is then used to mask all high-Z features out of the raw phase-contrast volume by replacing them with the theoretical low-Z μ value. This creates the image in (d), which is then filtered using 3D phase retrieval of the low-Z material [equation (2)] to make (e). Finally, the high-Z components from the two-material phase retrieved slice (b) are spliced into (e) using the binary mask, resulting in the final phase retrieved image (f), free of over-blurring artefacts from the high-Z boundaries. All non-binary images are displayed on the same colour palette which is optimized for viewing soft tissue features, leading to the bone in (f) appearing oversaturated and the subarachnoid space between the skull and brain tissue seeming indistinguishable from the surrounding agarose the object was set in.

interface (see further examples in Section 3) and potentially allows new features of the low- Z material to be resolved that would have previously been obscured by noise.

3. Results

To demonstrate the effectiveness of 3DMPR, we refer back to the data volume described in Section 2. By taking two cross-sections of the head CT volume in the sagittal plane, we compare the results of two phase retrieval approaches for the data. Figs. 3(a) and 3(d) show cross sections of a standard CT reconstruction from unfiltered propagation-based phase-contrast projections, where both slices display few-to-no discernible brain features. Following the standard practice of Croton *et al.* (2018a), we can construct a new CT volume from projections phase retrieved to the A/B material interface, namely brain-tissue–bone, using equation (3) [see Figs. 3(b) and 3(e)]. This allows resolution of some soft-tissue structures in the brain while others remain obscured by significant levels of noise still left in the reconstructions. Reducing the noise floor further requires stronger filtering, but merely applying

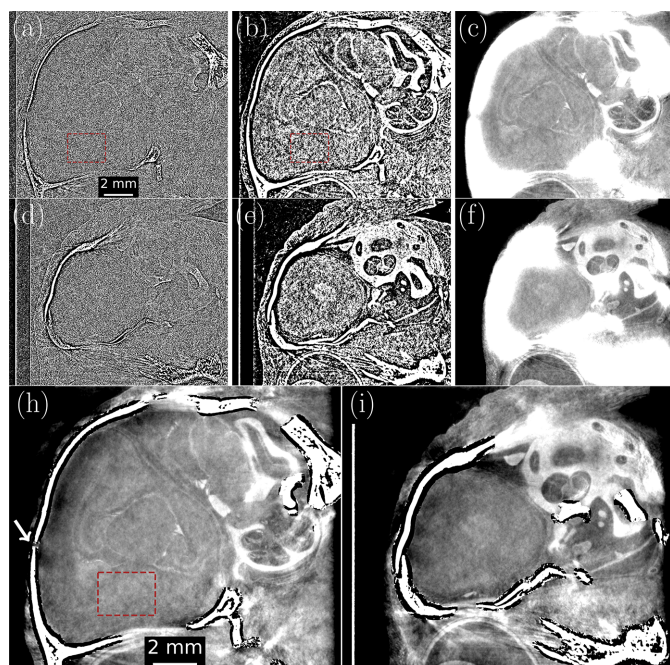


Figure 3

Comparison of our masked phase retrieval method with alternative approaches, shown using sagittal cross-sections of the rabbit head CT reconstruction volumes from Section 3. Panels (a) and (d) are two cross-sections of the phase contrast CT as a baseline for the other algorithms. Panels (b) and (e) are the bone–tissue interface phase retrieved image slices, equation (3), showing considerable remnant noise. Panels (c) and (f) show the result of applying equation (1) phase retrieval to each projection under the assumption that only brain tissue is present, leading to over-blurring completely obscuring brain tissue features in proximity to the bone. Panels (h) and (i) are retrieved using 3DMPR to achieve an equally low SNR image whilst correctly reconstructing the A/B material interface. Brown boxes denote regions used for SNR analysis while the white arrow in (h) highlights a gap in the animal skull that is not visible in the equation (1) phase retrieval. All images use the same colour palette with the region surrounding the animal filled with agarose.

phase retrieval for material A [equation (1)] in projection results in over-blurring of the high contrast interface, as seen in Figs. 3(c) and 3(f). This over-blurring leads to some features being obscured, for example the bone gap highlighted by the white arrow in Fig. 3(h) which is not visible using single material phase retrieval. These over-blurring artefacts may disrupt the linear attenuation coefficients far from the material boundaries, making large portions of the image no longer quantitative. The method presented in this paper, producing Figs. 3(h) and 3(i), achieves the noise reduction of stronger filtering while also retaining edge definition in the bone–brain boundary, avoiding contrast from the bone blurring into the adjacent brain tissue. Overall, we see features of the brain anatomy with greater clarity across the entire head than with the alternative reconstruction methods.

3.1. SNR characterization

SNR characterizations were performed using the regions indicated in Figs. 3(a), 3(b) and 3(h). Values were calculated in a five-slice radius, using the standard deviation as the error. Raw phase contrast reconstructions [Fig. 3(a)] produced a SNR of 1.09 ± 0.02 , which was improved under two-material phase retrieval [equation (3)] interface to 37.0 ± 0.3 . 3DMPR further increases the SNR by a factor of 6.8 to 252 ± 6 , representing a total SNR gain of 231, without compromising resolution. However, the region being analysed includes variations due to features of the brain tissue as shown in Fig. 3(h) which are recorded as noise measured through the standard deviation. The SNR gain of 6.8 is then likely an underestimation, but still represents a further potential dose reduction of $6.8^2 \simeq 46$ -fold (Kitchen *et al.*, 2017) that may become part of achieving lower-dose brain imaging.

3.2. Boundary preservation

To quantify the method's ability to preserve high contrast material boundaries and avoid over-blurring, we explore the simple and well characterized dataset of an aluminium pin submerged in water. This dataset was recorded on a $50 \mu\text{m}$ pixel Hamamatsu CMOS flat panel detector (C9728DK-10) using a microfocus X-ray source (THE-Plus from X-RAY WorX GmbH) with a silver transmission target at a tube voltage of 35 kV_p . The mean X-ray energy was determined by comparing μ values of water measured in CT at the depth of the aluminium pin, resulting in a mean energy of 19.58 keV . The source-to-object distance was 0.96 m with the source-to-detector distance set to 2.40 m , resulting in a $2.50\times$ magnification and effective propagation distance of 0.576 m . CTs were reconstructed with 3271 projections at 0.11° increments using filtered back-projection in a fan beam geometry (van Aarle *et al.*, 2015; van Aarle *et al.*, 2016). For this sample, material A is taken to be water, $\mu = 84.72 \text{ m}^{-1}$ and $\delta = 6.00 \times 10^{-7}$, while material B is the comparatively high- Z aluminium with parameters $\mu = 985.86 \text{ m}^{-1}$ and $\delta = 1.38 \times 10^{-6}$ (Schoonjans *et al.*, 2011).

Fig. 4 shows example portions of a CT slice through the aluminium pin for phase retrieval (b) with equation (1) and

(c) with our 3DMPR method. Each figure applies the same filtering to regions surrounding the aluminium cylinder, but (c) uses two-material phase retrieval to more accurately preserve the A/B material boundary. For 3DMPR, the aluminium cylinder was masked using a μ threshold of 300 m^{-1} with two iterations of a $3 \times 3 \times 3$ dilation filter. Despite the aluminium insert having uniform density, both Figs. 4(b) and 4(c) show a decrease in attenuation coefficient toward the centre, which is a cupping artefact caused by beam hardening (Barrett & Keat, 2004). This results from variable penetration depths of the polychromatic X-ray spectrum (Croton *et al.*, 2018a). Fig. 4(d) plots azimuthally averaged line profiles across the A/B boundary in panels (a), (b) and (c), incorporating the phase contrast profile from the non-phase retrieved CT slice for comparison. The solid vertical line represents the edge of the mask used for replacing material B. The horizontal dashed line denotes the theoretical linear attenuation coefficient value used for the temporary high-Z replacement, calculated using the mean energy of the polychromatic source, which closely resembles the measured linear attenuation coefficient outside the region being replaced. Fig. 4(d) shows the edge profile of our 3DMPR approach with dark grey datapoints. This more closely resembles the true boundary interface of the aluminium pin, whereas the equation (1) phase retrieval approach, shown in Fig. 4(d) with the black datapoints, produces a low-contrast blurred edge due to over-blurring, even when retrieving at a relatively small effective propagation distance of 0.576 m. This over-blurring leaves the area around the boundary unsuitable for quanti-

tative analysis, even for any features of interest not obscured by the edge blurring. To measure the blurring extent of each approach, we numerically differentiated each edge profile and applied a Pearson VII fit to determine the full width at half-maxima of the effective point spread function (PSF) (Croton *et al.*, 2018a). Pearson VII functions are appropriate fits for PSFs (Croton *et al.*, 2018a) since they have freedom that enables their shape to lie on a spectrum between Lorentzian and Gaussian. Through this measure, the spatial resolution of the conventional phase retrieval approach was found to be $108 \pm 2 \mu\text{m}$, while our approach was measured at $54 \pm 1 \mu\text{m}$. This demonstrates a clear improvement in material boundary definition while retaining the same SNR boost in the low-Z material of $4.2\times$, measured using the same method as described in Section 3.1. This SNR boost will be particularly valuable in improving weakly attenuating feature resolution in low flux imaging scenarios.

3.3. Expansion to 3+ material interfaces

While Section 2 demonstrated the 3DMPR process using a two-material sample, here we generalize to higher sample compositions and demonstrate how masked phase retrieval may be used to overcome a restriction of the two-material algorithm of equation (3). The phase retrieval algorithms of equations (1) and (3) require that each material is only in contact with a maximum of one other material or vacuum. Otherwise, only parts of the material's boundary will be quantitatively reconstructed.

Phantoms containing relevant information in multiple attenuation levels adjacent to high-Z boundaries were not imaged during experimentation; therefore, to demonstrate the concept we consider a mock three-material CT slice, shown in Fig. 5(a), composed of materials I, II and III and possessing four unique interfaces: I/air, I/II, I/III and II/III. Phase retrieval with the two-material algorithm [equation (3)] cannot fully resolve the boundary of any material. For example, phase retrieval focused on the II/I interface will leave the II/III interface non-quantitative. Our approach relaxes the requirement for spatially isolated materials but does require materials to be contrast-resolved, as in the histogram of Fig. 5(b). This allows two-sided thresholding to create binary masks for each material, Figs. 5(c), 5(d) and 5(e), in place of the single-sided thresholding used in Section 2. Although the material distributions in Fig. 5(b) may appear idealized, 3DMPR is proposed to counter over-blurring from phase retrieval specifically where high-Z materials are present in a comparably low-Z medium. Otherwise, materials with overlapping μ distributions may share equivalent enough phase characteristics to be considered a single material, allowing an average value to be used in conventional TIE phase retrieval and negating the need for 3DMPR. Alternatively, other mask segmentation methods may be employed, such as region-growing algorithms or other advanced processes that utilize position information and that would benefit from phase-contrast fringes highlighting sample boundaries, particularly in the presence of noise.

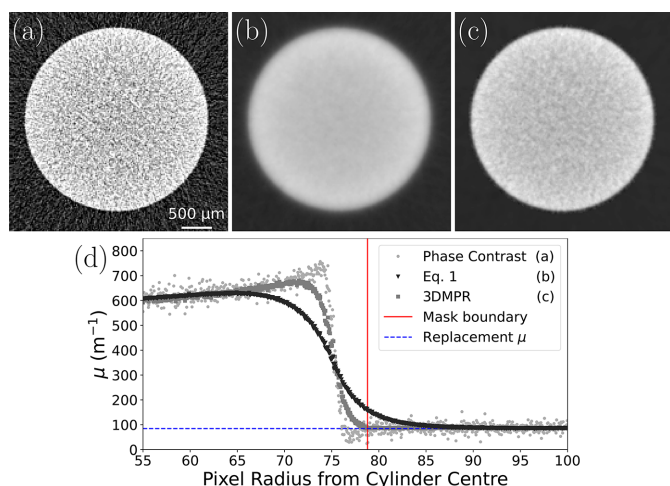


Figure 4 Quantification of edge resolution in a 3 mm-thick aluminium rod submerged in water after phase retrieval was applied to material A, water. Panel (a) shows a phase contrast slice of the aluminium cylinder while panels (b) and (c) show the same slice after from phase retrievals performed through the single-material algorithm [equation (1)] and 3DMPR, respectively. Panel (d) plots line profiles produced by azimuthally averaging around the entire aluminium–water interface to compare the raw phase contrast resolution with that of the conventional 2D phase retrieval and our 3D masking method. A solid vertical line is used to show the mask radius, notably not visible in (c), while a horizontal line shows the theoretical μ value of the low-Z material (water) used for replacing the masked aluminium in the 3D reconstruction. Images (b) and (c) share the same colour palette.

3DMPR is then divided into two categories: the relatively constant regions within each material, termed homogeneous regions, and regions around material boundaries containing two materials, referred to as heterogeneous. Phase retrieval for homogeneous regions is performed to the CT volume, using equation (1), after replacing all other image components with the theoretical material value using the inverse of each mask [Figs. 5(c), 5(d) and 5(e)]. A dilation filter may be required on the inverse masks to ensure that heterogeneous regions are not included. When combined, these reconstructions provide strong noise suppression within each material,

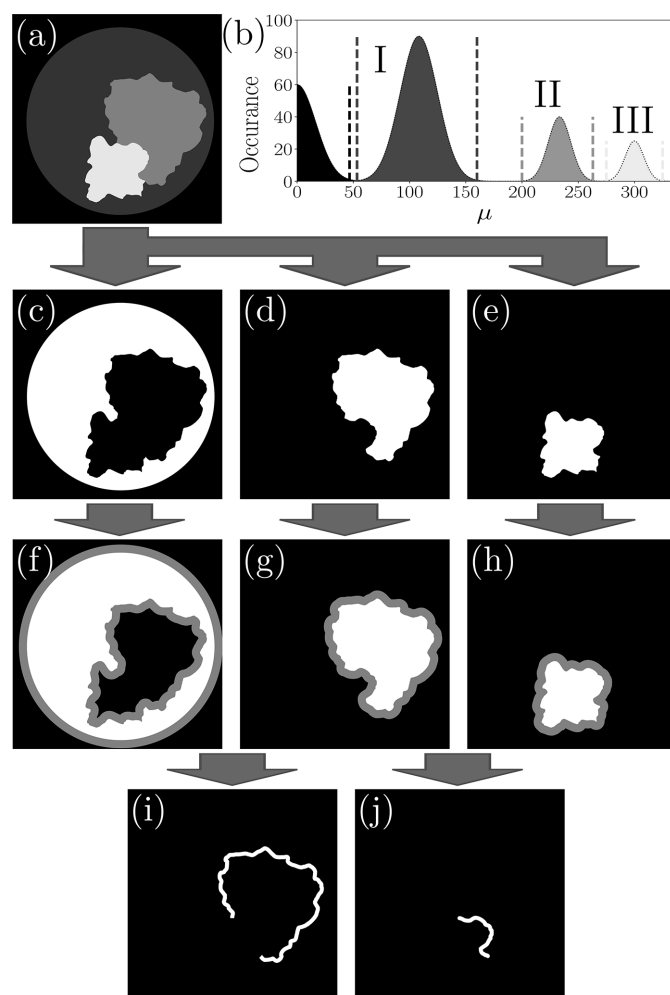


Figure 5 Conceptualized application of 3DMPR to a multi-material object composed of materials I, II and III, as shown in (a). Images shown represent single slices of the CT volume with the assumption of contrast isolation between all materials as in the histogram (b). Dashed, vertical lines demonstrate example threshold bounds for creating material masks, shown in (c), (d) and (e). The inverse of these masks, black regions, can be used to mask away all other material components, allowing phase retrieval to be performed to internal material regions using equation (1). The boundary regions are interleaved from appropriately tuned retrievals, located using binary operations between the material masks after an additional dilation (f), (g), (h). Panel (i) shows the result of a binary ‘and’ operation between the material I and II masks, creating a trace of the I/II interface for interleaving. Panel (j) demonstrates the same for the II/III interface, overcoming the convention limit of two-way material boundaries.

without over-blurring low- Z components, but do not yet include the material boundaries. These heterogeneous sample regions must be interleaved from appropriately tuned reconstructions, requiring knowledge of the boundary locations, acquired through further manipulation of the binary masks [Figs. 5(c), 5(d), 5(e)]. The masks are expanded by an equivalent amount via dilation filters [Figs. 5(f), 5(g), 5(h)] until masks of adjacent materials overlap. Binary ‘and’ operations are then used to isolate the overlapping regions which will be centred about the material boundaries. Figs. 5(i) and 5(j) show examples of this for the II/I and II/III material interfaces, respectively, which can then be used to select these regions from the appropriate reconstructions.

Interleaving all homogeneous and heterogeneous regions together, and iterating the method proven in Sections 2, 3.1 and 3.2, will produce a noise suppressed CT reconstruction with quantitative resolution at all material boundaries.

4. Discussion

3DMPR is a method for interface-specific phase retrieval which overcomes filtering limitations at the boundaries of high- Z materials while maintaining the same assumptions as equation (1), *i.e.* prior knowledge of the sample composition and resolution of phase contrast fringes. Note that, if the materials are not known, their properties can nevertheless be determined by analysing the phase contrast fringes in the CT (Alloo *et al.*, 2022). Although clear material distinctions in CT are required to create the 3D mask, this is a defining aspect of the posed problem of applying phase retrieval near highly attenuating objects. If such material distinctions are not present then regular phase retrieval may be applied instead. Applying 3DMPR provides greater high frequency noise suppression than equation (3) alone, increasing the clarity of all features including any artefacts already present in the CT reconstruction. This is shown in Fig. 2 where streak artefacts (Barrett & Keat, 2004) are visible from photon trajectories parallel to the edge of the bones. Because the streak artefacts arise during the CT reconstruction step, and not the phase retrieval step, the approach described in this paper will not address these artefacts. Similar to the Paganin and Beltran algorithms, 3DMPR will provide the greatest benefits for large propagation distances, low- Z materials and low energies, subject to conditions satisfying the TIE. This includes requiring sufficiently small pixel sizes to resolve the phase contrast fringes. Specifically, if $(\delta_2 - \delta_1)/(\mu_2 - \mu_1)$ is smaller than δ/μ used in 3DMPR then the SNR will improve.

We believe 3DMPR will increase the application scope of phase retrieval, by improving results in the presence of highly attenuating materials, such as the medical bone–flesh examples we present in Figs. 2 and 3. More complex extensions with three materials could add a distinction for the cartilage seen in Figs. 3(b) and 3(e). Similar medical applications could include distinguishing soft material from highly attenuating implants/contrast agents in and around bone where three materials may meet, creating a scenario where standard phase retrieval has difficulties and neutron imaging may otherwise be required

(Isaksson *et al.*, 2017). Other examples of complex sample applications could be analysing soil samples or manufactured objects, in particular with metals or minerals. Each phantom would require an initial optimization of the masking process. For example, see the difference in dilation filters between the datasets of Figs. 3 and 4 where the former required many more dilations to cover the low attenuation region between the skull and brain tissue, known as the subarachnoid space, when starting from the bone threshold as an initial seed. However, the processing of similar samples would follow the same settings, which could be used to establish an efficient workflow with no further manual interventions required.

The problem of phase retrieval with two-material samples has also been tackled before using successive retrieval on the one volume (Ullherr & Zabler, 2015), including presenting what we believe is the first 3D phase retrieval filter and a similar method to ours for retrieving near high-*Z* materials (Ullherr & Zabler, 2015). The primary differences between each method are that Ullherr & Zabler (2015) apply the second phase retrieval directly to the equation (3) phase retrieved CT and that they do not perform material replacement in the masked regions. Applying 3DMPR to the two-material retrieved CT volume, instead of a phase contrast volume, may provide a slightly higher SNR due to phase retrieval being more noise tolerant than phase contrast CT. However, replacing the mask area creates an effectively homogeneous material, eliminating over-blurring from neighbouring regions or erosion of features due to neighbouring masked regions without the need for additional post-processing. Haggmark *et al.* (2017) further explore the Ullherr & Zabler approach (Ullherr & Zabler, 2015) alongside two-material phase retrieval (Beltran *et al.*, 2011) using the labels ‘linear’ and ‘parallel’, respectively.

5. Conclusion

We present a method for achieving strong noise reduction during phase retrieval of multi-material objects, without creating over-blurring along material boundaries with high-*Z* materials. By using simple thresholding techniques for material isolation in CT, the method is computationally efficient and compatible with polychromatic reconstructions. A rabbit kitten brain CT was used to demonstrate the algorithm and its ability to improve image quality in complex biological samples. SNR values increased from 1.08 ± 0.02 in the phase contrast images to 252 ± 6 using 3DMPR, an increase of 6.8 times over the two-material phase retrieval algorithm, whilst preserving the same boundary definition between the high- and low-*Z* materials. Conversely, when applied to a separate, well characterized dataset, edge resolution across the high-*Z* boundary was markedly improved from $108 \pm 2 \mu\text{m}$ using the approach in equation (1) to $54 \pm 1 \mu\text{m}$ with the new approach. This method will allow strong SNR-boosting phase retrieval to be applied to a wider range of multi-material, including those with three or more material interfaces.

Acknowledgements

This experiment used rabbit kittens that had been used in experiments conducted with approval from the SPring-8 Animal Care (Japan) and Monash University (Australia) Animal Ethics Committees. Open access publishing facilitated by Monash University, as part of the Wiley–Monash University agreement via the Council of Australian University Librarians.

Data availability

The datasets and/or analysis pipelines used for this manuscript are available from the corresponding author on reasonable request.

Funding information

JAP is supported by a Research Training Program (RTP) Scholarship and the J. L. Williams Top-up Scholarship. KM acknowledges support from the Australian Research Council (FT18010037). SBH is an NHMRC Principal Research Fellow. This work was funded by NHMRC 2021 Ideas Grants Application 2012257 and supported by the Victorian Government’s Operational Infrastructure Support Program.

References

- Albers, J., Wagner, W. L., Fiedler, M. O., Rothermel, A., Wünnemann, F., Di Lillo, F., Dreossi, D., Sodini, N., Baratella, E., Confalonieri, M., Arfelli, F., Kalenka, A., Lotz, J., Biederer, J., Wielpütz, M. O., Kauczor, H., Alves, F., Tromba, G. & Dullin, C. (2023). *Sci. Rep.* **13**, 4788.
- Alloo, S. J., Paganin, D. M., Morgan, K. S., Gureyev, T. E., Mayo, S. C., Mohammadi, S., Lockie, D., Menk, R. H., Arfelli, F., Zanconati, F., Tromba, G. & Pavlov, K. M. (2022). *Opt. Lett.* **47**, 1945.
- Arhatari, B. D., Nesterets, Y. I., Taba, S. T., Maksimenko, A., Hall, C. J., Stevenson, A. W., Hausermann, D., Lewis, S. J., Dimmock, M., Thompson, D., Mayo, S. C., Quiney, H. M., Gureyev, T. E. & Brennan, P. C. (2021). *Proc. SPIE* **11840**, 1184012.
- Barrett, J. F. & Keat, N. (2004). *Radiographics* **24**, 1679–1691.
- Beltran, M. A., Paganin, D. M., Siu, K. K. W., Fouras, A., Hooper, S. B., Reser, D. H. & Kitchen, M. J. (2011). *Phys. Med. Biol.* **56**, 7353–7369.
- Beltran, M. A., Paganin, D. M., Uesugi, K. & Kitchen, M. J. (2010). *Opt. Express* **18**, 6423.
- Berger, M. J., Hubbell, J. H., Seltzer, S. M., Chang, J., Coursey, J. S., Sukumar, R., Zucker, D. S. & Olsen, K. (2010). *XCOM: Photon cross sections database NIST Standard Reference Database 8 (XGAM)*. National Institute of Standards and Technology, Gaithersburg, USA.
- Bidola, P., Morgan, K., Willner, M., Fehringner, A., Allner, S., Prade, F., Pfeiffer, F. & Achterhold, K. (2017). *J. Microsc.* **266**, 211–220.
- Brody, D. L., Mac Donald, C. L. & Shimony, J. S. (2015). *Handbook of Clinical Neurology* edited by J. Grafman & A. M. Salazar, Volume 127 of *Traumatic Brain Injury, Part I*, ch. 17, pp. 267–275. Elsevier.
- Chantler, C. T., Olsen, K., Dragoset, R. A., Chang, J., Kishore, A. R., Kotochigova, S. A. & Zucker, D. S. (2005). *X-ray Form Factor, Attenuation, and Scattering Tables NIST Standard Reference Database 66*. National Institute of Standards and Technology, Gaithersburg, USA.

- Croton, L., Morgan, K., Paganin, D., Kerr, L., Wallace, M., Crossley, K., Ruben, G., Miller, S., Yagi, N., Uesugi, K., Hooper, S. & Kitchen, M. (2018b). *Microsc. Microanal.* **24**, 354–355.
- Croton, L. C. P., Morgan, K. S., Paganin, D. M., Kerr, L. T., Wallace, M. J., Crossley, K. J., Miller, S. L., Yagi, N., Uesugi, K., Hooper, S. B. & Kitchen, M. J. (2018a). *Sci. Rep.* **8**, 11412.
- Dullin, C., dal Monego, S., Larsson, E., Mohammadi, S., Krenkel, M., Garrovo, C., Biffi, S., Lorenzon, A., Markus, A., Napp, J., Salditt, T., Accardo, A., Alves, F. & Tromba, G. (2015). *J. Synchrotron Rad.* **22**, 1465–1455.
- Gureyev, T. E., Mayo, S. C., Nesterets, Y. I., Mohammadi, S., Lockie, D., Menk, R. H., Arfelli, F., Pavlov, K. M., Kitchen, M. J., Zanconati, F., Dullin, C. & Tromba, G. (2014). *J. Phys. D Appl. Phys.* **47**, 365401.
- Gureyev, T. E., Nesterets, Y. I., Kozlov, A., Paganin, D. M. & Quiney, H. M. (2017). *J. Opt. Soc. Am. A* **34**, 2251.
- Häggmark, I., Vågberg, W., Hertz, H. M. & Burvall, A. (2017). *Opt. Express* **25**, 33543.
- Hehn, L., Gradl, R., Voss, A., Günther, B., Dierolf, M., Jud, C., Willer, K., Allner, S., Hammel, J. U., Hessler, R., Morgan, K. S., Herzen, J., Hemmert, W. & Pfeiffer, F. (2018b). *Biomed. Opt. Expr.* **9**, 5330.
- Hehn, L., Morgan, K., Bidola, P., Noichl, W., Gradl, R., Dierolf, M., Noël, P. B. & Pfeiffer, F. (2018a). *APL Bioeng.* **2**, 016105.
- Isaksson, H., Le Cann, S., Perdikouri, C., Turunen, M. J., Kaestner, A., Tägil, M., Hall, S. A. & Tudisco, E. (2017). *Bone* **103**, 295–301.
- Kitchen, M. J., Buckley, G. A., Gureyev, T. E., Wallace, M. J., Andres-Thio, N., Uesugi, K., Yagi, N. & Hooper, S. B. (2017). *Sci. Rep.* **7**, 15953.
- Lang, J. A. R., Pearson, J. T., Binder-Heschl, C., Wallace, M. J., Siew, M. L., Kitchen, M. J., te Pas, A. B., Fouras, A., Lewis, R. A., Polglase, G. R., Shirai, M. & Hooper, S. B. (2016). *J. Physiol.* **594**, 1389–1398.
- Lang, J. A. R., Pearson, J. T., te Pas, A. B., Wallace, M. J., Siew, M. L., Kitchen, M. J., Fouras, A., Lewis, R. A., Wheeler, K. I., Polglase, G. R., Shirai, M., Sonobe, T. & Hooper, S. B. (2014). *J. Appl. Physiol.* **117**, 535–543.
- Lioliou, G., Buchanan, I., Astolfo, A., Endrizzi, M., Bate, D., Hagen, C. K. & Olivo, A. (2024). *Opt. Express* **32**, 4839.
- Mohammadi, S., Larsson, E., Alves, F., Dal Monego, S., Biffi, S., Garrovo, C., Lorenzon, A., Tromba, G. & Dullin, C. (2014). *J. Synchrotron Rad.* **21**, 784–789.
- Nesterets, Y. I., Gureyev, T. E., Mayo, S. C., Stevenson, A. W., Thompson, D., Brown, J. M. C., Kitchen, M. J., Pavlov, K. M., Lockie, D., Brun, F. & Tromba, G. (2015). *J. Synchrotron Rad.* **22**, 1509–1523.
- Paganin, D., Mayo, S. C., Gureyev, T. E., Miller, P. R. & Wilkins, S. W. (2002). *J. Microsc.* **206**, 33–40.
- Pelliccia, D., Kitchen, M. J. & Morgan, K. S. (2018). *Handbook of X-ray Imaging* pp. 971–998. CRC Press.
- Pollock, J. A., Morgan, K., Croton, L. C. P., Pryor, E. J., Crossley, K. J., Hall, C. J., Häusermann, D., Maksimenko, A., Hooper, S. B. & Kitchen, M. J. (2025). *Sci. Rep.* **15**, 23546.
- Reichmann, J., Ruhwedel, T., Möbius, W. & Salditt, T. (2023). *J. Med. Imag.* **10**, 056001.
- Schoonjans, T., Brunetti, A., Golosio, B., Sanchez del Rio, M., Solé, V. A., Ferrero, C. & Vincze, L. (2011). *At. Spectrosc.* **66**, 776–784.
- Temple, N., Donald, C., Skora, A. & Reed, W. (2015). *J. Med. Radiat. Sci.* **62**, 122–131.
- Thompson, D. A., Nesterets, Y. I., Pavlov, K. M. & Gureyev, T. E. (2019). *J. Synchrotron Rad.* **26**, 825–838.
- Ullherr, M. & Zabler, S. (2015). *Opt. Express* **23**, 32718.
- van Aarle, W., Palenstijn, W. J., Cant, J., Janssens, E., Bleichrodt, F., Dabrovolski, A., De Beenhouwer, J., Joost Batenburg, K. & Sijbers, J. (2016). *Opt. Express* **24**, 25129.
- van Aarle, W., Palenstijn, W. J., De Beenhouwer, J., Altantzis, T., Bals, S., Batenburg, K. J. & Sijbers, J. (2015). *Ultramicroscopy* **157**, 35–47.
- Vliegenthart, R., Fouras, A., Jacobs, C. & Papanikolaou, N. (2022). *Respirology* **27**, 818–833.