



Criteria for selecting the Paganin-filter reconstruction parameter in X-ray phase-contrast tomography

Eduardo X. Miqueles,^{a*} Everton L. de Oliveira,^a Murilo Carvalho^a and David M. Paganin^b

Received 13 October 2025
Accepted 17 February 2026

^aBrazilian Synchrotron Light Laboratory, CNPEM, Campinas, SP, Brazil, and ^bSchool of Physics and Astronomy, Monash University, Clayton, Australia. *Correspondence e-mail: eduardo.miqueles@lnls.br


Edited by A. Stevenson, Australian Synchrotron, Australia

Keywords: tomography; Paganin; noise; algorithm; imaging.

The Paganin filter, widely employed in propagation-based phase-contrast X-ray computed tomography with synchrotron light, attenuates phase effects and suppresses high-frequency noise under the assumption that samples have a constant β/δ ratio (these constants define the material's complex refractive index $n = 1 - \delta + i\beta$). While effective in many scenarios, real samples often violate this strong assumption, and the choice of the filter's free reconstruction parameter is typically made empirically, potentially leading to suboptimal image quality. This work introduces a quantitative, noise-aware criterion for determining a near-optimal filter parameter, thereby removing the need for subjective tuning. We propose two methods for this task. The first method is based on analysis of the noise power spectrum in radial frequency bands, adopting the same partitioning framework as Fourier ring correlation (FRC). By monitoring high-frequency behavior, the proposed approach identifies the trade-off point where resolution enhancement is balanced against noise amplification. This criterion enforces a theoretical lower bound on the FRC computed between two independently split Paganin-filtered projections, where each projection corresponds to a single recorded detector frame, while ensuring effective suppression of high-frequency noise. The second method analyzes the variation of the reconstructed images as a function of the Paganin-filter parameter, identifying the inflection point that represents the balance between noise suppression and image smoothing. Both methods result in an automated parameter selection approach for phase-contrast tomography data analysis that employs the Paganin filter.

1. Introduction

In computational methodologies applied to computed X-ray tomography, especially when performed with a synchrotron light source, it is common for propagation-based phase-contrast measurements to have their phase effects processed through application of the so-called Paganin filter (Teague, 1983; Paganin *et al.*, 2002). This filter is based on the assumption that the ratio between the imaginary part (β) and the real decrement (δ) of the complex refractive index $n = 1 - \delta + i\beta$ is approximately constant across the object. In many experiments, this hypothesis is considered as approximately valid, making the method widely adopted for reconstructing images with adequate quality. However, in experimental practice and under real data acquisition conditions, the sample under investigation will often not be accurately described as having a constant β/δ ratio. In particular, it is common to observe a high heterogeneity of compositions and structural variations, which do not strictly conform to the assumptions of the original 'single material' model. Nevertheless, the Paganin filter has become one of the most widely used approaches in



synchrotron-based X-ray phase contrast imaging, due in part to its ease of implementation, low computational cost, and ability to produce stable results even under challenging conditions (Gureyev *et al.*, 2014a; Nesterets & Gureyev, 2014; Kitchen *et al.*, 2017; Gureyev *et al.*, 2017; Paganin *et al.*, 2020). Previous studies have addressed the optimization of key parameters in propagation-based phase-contrast imaging, including the regularization parameter in Paganin's method (Beltran *et al.*, 2018) and the propagation distance in free-space imaging setups (Gureyev *et al.*, 2004), highlighting the importance of parameter tuning for accurate phase retrieval.

In addition to attenuating phase artifacts, the filter can reduce part of the noise present in the projections while still maintaining a satisfactory level of fidelity to the original data (Miqueles *et al.*, 2025a). This feature makes it particularly attractive for quantitative and qualitative analyses in high-resolution experiments. However, in order for the result to be truly satisfactory, it is essential to appropriately determine the free parameter that is associated with the filter, which controls the degree of smoothing applied. In practice, this choice is often made empirically, being left to the experimental station user, who adopts an *ad hoc* value based on prior experience or preliminary testing.

To eliminate the need for manually choosing a value for the Paganin filter parameter ℓ , there is a clear demand for the development of computational strategies capable of providing the user with a near-optimal value, avoiding reliance solely on subjective judgment or trial-and-error. The goal is to offer a systematic procedure, based on quantitative metrics, that not only automates parameter selection but also reduces the uncertainty associated with this critical processing step. We emphasize that automatic methodologies for finding refractive properties of the material under investigation were already reported in the literature (Eastwood *et al.*, 2011; Alloo *et al.*, 2022), although from quite different computational perspectives in comparison with that reported in the present paper. The near-optimal value for the filter parameter, associated with our work, should satisfy different practical requirements simultaneously. On one hand, it should allow a clear separation of the histogram peaks of the obtained projection; this is a criterion that is often used visually, by experienced users, to assess the quality of contrast between regions of the sample. On the other hand, it should be capable of preserving or even enhancing edges and structural features of the object under analysis, thus maximizing the usefulness of the reconstruction for both quantitative and qualitative interpretations.

In this context, the present manuscript addresses the development of a criterion for near-optimal selection of the Paganin-filter parameter ℓ , based on noise power analysis measured via the image's noise power spectrum (NPS) (Nesterets & Gureyev, 2014). The proposed approach aligns with methodologies widely used for spatial resolution assessment, such as Fourier ring correlation (FRC) (Saxton & Baumeister, 1982; Banterle *et al.*, 2013; Nieuwenhuizen *et al.*, 2013; Culley *et al.*, 2018; Tortarolo *et al.*, 2018) and its three-dimensional variant, Fourier shell correlation (FSC). The central idea is to measure the NPS of the input signal over

radial frequency bands, applying a ring partitioning analogous to that used in FRC. This approach is compared with a second method, also proposed in this work, which is based on the standard deviation of the derivative of images with respect to the free parameter ℓ .

This approach offers two main advantages. First, it enables refined monitoring of high-frequency behavior, making it possible to identify the point at which the gain in sharpness is counterbalanced by undesired amplification of noise. Second, it provides a quantitative criterion that naturally aligns with the classical trade-off between noise reduction and fidelity to the original data. Thus, the proposed method not only systematizes parameter selection but also contributes to greater rigor in the analysis and comparison of results obtained in X-ray phase-contrast tomography experiments.

Since the approach developed in this paper may be viewed in the broader context of autofocusing computational imaging systems (Eastwood *et al.*, 2011), we close this introduction with some general remarks in this regard. Tomography, in general, and X-ray phase-contrast tomography, in particular, is an exemplar of the computational-imaging concept in the sense that the computer forms an intrinsic part of the imaging system, a 'software lens' which utilizes a mathematical-physics model of the forward-imaging problem in order to solve the inverse problem of tomographic reconstruction from a set of two-dimensional projection images (Paganin *et al.*, 2004). At no stage is the value at an individual voxel (*e.g.* a density) directly measured by a pixellated detector, thus the tomographic reconstruction may indeed be viewed as a 'software lens'. As is the case with the 'hardware lens' in the human visual imaging system, the 'software lens' associated with Paganin-filtered phase-contrast tomography depends on a single parameter; this single parameter corresponds to a defocus for the former (hardware) imaging system, and a β/δ ratio for the latter (computational/software) imaging system. In analogy to the autofocusing property of the human visual system—whereby we adaptively select an accurate focal length for our eyes' lenses in order to see an acceptably sharp image—the present paper may be viewed as a computational-imaging procedure whose associated 'software lens' (namely Paganin-method phase retrieval followed by tomographic reconstruction) is 'autofocused' in the systematic manner referred to earlier in this introduction.

2. Methodology

We briefly recall the action of the Paganin filter, which depends on a real parameter ℓ (Paganin *et al.*, 2002). In the special case of the transport-of-intensity equation (TIE) (Teague, 1983) that corresponds to a sample composed of a single homogeneous material of possibly variable density, the projection approximation implies that a propagation-based X-ray phase contrast intensity image f is mapped into another image p_ℓ through application of a convolution filter P_ℓ , followed by the logarithm applied to valid pixels, namely

$$f \mapsto p_\ell = -\log P_\ell(f), \quad P_\ell = K_\ell \star f, \quad (1)$$

with

$$\hat{K}_\ell(\mathbf{q}) = \frac{1}{1 + \ell \|\mathbf{q}\|^2}, \quad \mathbf{q} \in \mathbb{R}^2.$$

Here the \star symbol denotes two-dimensional convolution, $\hat{\cdot}$ denotes two-dimensional Fourier transformation with respect to Cartesian spatial coordinates (x, y) , $\mathbf{q} \equiv (q_x, q_y)$ are the corresponding Fourier coordinates (spatial-frequency coordinates) having units of inverse length, $\|\mathbf{q}\|^2 = q_x^2 + q_y^2$, and the filter parameter ℓ has units of squared length. The convolution kernel acts as a low-pass filter in the spatial-frequency domain, attenuating high-frequency components—often associated with noise—while preserving the main structural features of the image. Under the previously stated assumptions, together with the requirements of uniform-intensity coherent plane-wave illumination and a sample-to-detector distance which is sufficiently small for the associated Fresnel number to be much greater than unity, p_ℓ is approximately proportional to the projected thickness of the homogeneous sample, with this projection being taken along the direction of the optical axis.

It is important to note a distinction with respect to filter-parameter values more commonly used in the literature, where the parameter is often expressed as the dimensionless quotient $\ell_L = \beta/\delta$, with δ and β representing the real and imaginary components, respectively, of the sample's complex refractive index $n = 1 - \delta + i\beta$ (Paganin *et al.*, 2002). In our formulation, we instead adopt the Paganin-filter parameter

$$\ell = z \frac{\delta}{\mu},$$

where z is the sample-to-detector propagation distance and $\mu = 4\pi\beta/\lambda$ is the linear attenuation coefficient, with λ denoting the X-ray wavelength. This grouping makes explicit the dependence of ℓ on experimental geometry and beam energy. From this definition, it follows directly that

$$\ell = \frac{z\lambda}{4\pi\ell_L}, \quad (2)$$

which ensures dimensional consistency of the convolution kernel in equation (1) when expressed in Fourier space. Thus, the uncertainties typically associated with the ratio β/δ are inherently transferred to ℓ through equation (2). In other words, any imprecision in the refractive index-parameters is reflected proportionally in the effective filter parameter ℓ . Once ℓ is determined experimentally or through prior sample-material knowledge, the conversion in equation (2) allows computation of ℓ_L . For completeness, we recall that in practice the input image f is the ratio between two other images typically acquired in synchrotron-based X-ray phase contrast tomographic experiments, namely $f = I/I_0$, where I is the photon count at the detector with the beam passing through the sample and I_0 is the corresponding flat-field image acquired without the sample in the beam path. This normalization step removes the incident beam profile from the data, ensuring that the Paganin-method filtering operates on a corrected intensity distribution representative of the sample alone.

2.1. Noise power spectrum criteria

When analyzing the spatial resolution of a single image, it is common practice to split it into two subsets (*e.g.* pixels with even and odd indices) so that the resulting images can be cross-correlated—for instance, via FRC—in order to estimate the resolution limit. In this work, the splitting procedure preserves the statistical nature of the original image, as demonstrated by Rieger *et al.* (2024), making it valid, for example, for resolution analysis. If we denote the original image by f , its split into two complementary images $\{u, v\}$ still contains sufficient information regarding the noise structure. This type of image splitting can be implemented in a practical way (Banterle *et al.*, 2013) with the purpose of determining resolution in microscopy.

The NPS (Nesterets & Gureyev, 2014) of a noise signal n is estimated as the function

$$W(\mathbf{q}) = \mathbb{E} \left[|\hat{n}(\mathbf{q})|^2 \right], \quad \mathbf{q} \in \mathbb{R}^2, \quad (3)$$

where $\mathbb{E}[\cdot]$ denotes the mathematical expectation of a random variable. When the original image is split into u and v , the noise content is preserved in each subset, so that the difference $n = u - v$ retains the same statistical noise structure as the original signal.

It is important to emphasize that equation (3) is employed here in the operational sense customary in imaging science. The objects under consideration are finite, discretely sampled detector frames with compact support, rather than idealized infinite stationary random processes. Accordingly, the NPS is estimated via the discrete Fourier transform of a finite noise realization (obtained in this work from the split-image difference), using a periodogram-type estimator. The expectation operator is approximated by averaging over Fourier samples within radial frequency bands, under the usual assumptions of local stationarity and isotropy across the finite field of view. In this discrete finite-dimensional setting, the NPS corresponds to the standard power spectral density (PSD) estimator used in detector and imaging-system characterization [*cf.* Goodman (2000); Dainty & Shaw (1974); Nesterets & Gureyev (2014)], and should not be interpreted as the Fourier transform of an idealized infinite stationary random process.

In order to avoid explicit computation of the expectation operator, we approximate $\mathbb{E}[\cdot]$ by an average over spatial frequencies, grouping together Fourier coefficients that belong to equivalent frequency domains, under the assumptions of stationarity and isotropy. These equivalent frequency domains are defined as *radial frequency bands*, in exactly the same manner as the rings used in FRC calculations. Here, the radial variable r denotes a normalized radial frequency coordinate. More precisely, if $\rho = \|\mathbf{q}\|$ represents the physical radial spatial frequency (with units of inverse length), and ρ_{\max} denotes the maximum sampled frequency in the discrete Fourier grid (*i.e.* the Nyquist radius), we define

$$r = \frac{\rho}{\rho_{\max}} \in [0, 1].$$

Thus, throughout this work, r is dimensionless and represents the relative position within the Nyquist band. In this sense, statements such as $0 < r < 1$ and choices like $r \simeq 0.5$ refer to fractions of the maximum resolvable spatial frequency (e.g. half the Nyquist radius) rather than to absolute spatial frequencies expressed in physical units.

Thus, the NPS can be expressed in terms of frequency bands of Fourier-space radius r as

$$W(r) = \frac{1}{2A |S_r|} \sum_{\mathbf{q} \in S_r} |\hat{u}(\mathbf{q}) - \hat{v}(\mathbf{q})|^2, \quad (4)$$

where S_r is the set of Fourier samples (pixels) in the ring corresponding to frequency r , $|S_r|$ is the number of elements in that set, and A is the total physical area of the image in real space. This normalization ensures that $W(r)$ is expressed in consistent physical units (e.g. intensity² × mm²) and can be directly compared across different acquisition conditions or detector systems.

Criterion. If we take an image p , obtained from the Paganin filter with parameter ℓ applied to the original signal f , we can state that its split into u_ℓ and v_ℓ is such that, as ℓ increases indefinitely, u_ℓ and v_ℓ become progressively more similar. This claim follows directly from the definition in equation (1), using the continuity of the logarithmic function, and can be verified with minimal computation. In geometric terms, as ℓ grows, the angle between the images—considered as vectors in a suitable function space (Barrett & Myers, 2004)—tends toward zero. In physical terms, increasing ℓ increases the strength of the low-pass filtration of the input propagation-based phase contrast image, thereby giving progressively higher relative weight to the (noise insensitive) lowest-order Fourier frequencies in the power spectrum of the measured signal. Naturally, this implies that the correlation of u_ℓ and v_ℓ approaches 1 in a monotonic fashion.

Mathematically, this behavior is made explicit through the FRC between u_ℓ and v_ℓ , expressed via the curve c_ℓ as a function of the radial frequency band r in Fourier space, namely,

$$c_\ell(r) = \frac{\langle \hat{u}_\ell(r), \hat{v}_\ell(r) \rangle}{\|\hat{u}_\ell(r)\| \|\hat{v}_\ell(r)\|}, \quad (5)$$

where ℓ is fixed for the computation of $c_\ell(r)$ but may vary when comparing results for different filter strengths. Above, $\|\cdot\|$ denotes a norm induced by an inner product $\langle \cdot, \cdot \rangle$ in the function space L_2 (Barrett & Myers, 2004).

It was demonstrated recently (Miqueles *et al.*, 2025b) that, for any fixed ℓ , the curve $c_\ell(r)$ satisfies a theoretical lower bound $B_\ell(r)$ (see Miqueles *et al.*, 2025b), which depends parametrically on ℓ ,

$$B_\ell(r) = \frac{2}{[a_\ell(r)]^{1/2} + [a_\ell(r)]^{-1/2}}, \quad a_\ell(r) = M_\ell(r)/m_\ell(r), \quad (6)$$

where $M_\ell(r)$ and $m_\ell(r)$ denote, respectively, the maximum and minimum magnitudes of the Fourier coefficients of the split images over the ring S_r . In other words, for all admissible radial frequencies r , one has $c_\ell(r) \geq B_\ell(r)$. As ℓ increases,

the strength of the low-pass filtering increases, leading to progressively higher correlation between the split images across all frequency bands. Consequently, for each fixed r , the value $c_\ell(r)$ increases and tends toward 1, corresponding to perfect correlation (i.e. the horizontal line $c = 1$ in the limit of very strong filtering). This qualitative trend is illustrated schematically in Fig. 1(a).

From this perspective, u_ℓ and v_ℓ behave as function-space vectors whose mutual angle decreases toward zero in the limit $\ell \rightarrow \infty$, making them nearly linearly dependent. The trade-off we seek is to determine the point at which we should stop increasing ℓ , so that this angle does not vanish entirely. This is where the NPS-based criterion comes into play, providing a quantitative way to balance how much this angle is allowed to close. Proposition 1 summarizes the algorithmic methodology we aim to adopt.

Proposition 1. Let $c_\ell(r)$ denote the split-image FRC curve defined in equation (5), and let $B_\ell(r)$ be the corresponding consistency lower bound given in equation (6). Assume that, for each fixed ℓ , the consistency inequality

$$c_\ell(r) \geq B_\ell(r) \quad \forall r \in (0, 1]$$

holds (cf. Miqueles *et al.*, 2025b). Define $\ell_{r_{\min}}$ as the smallest filter parameter ℓ such that the normalized high-frequency noise level satisfies

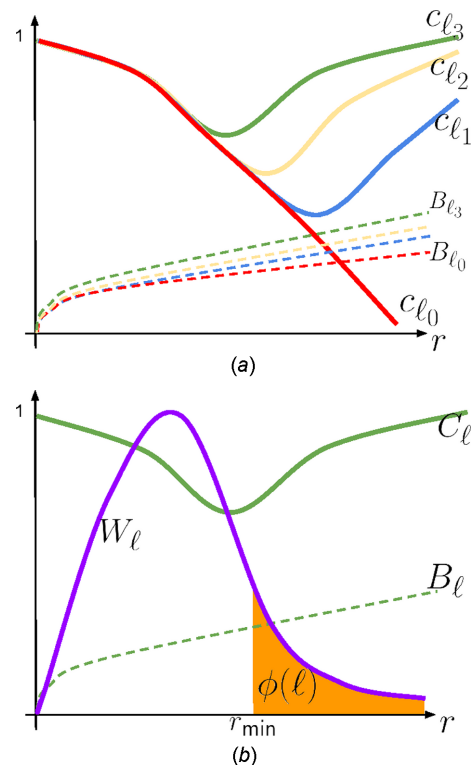


Figure 1 (a) Typical shape of correlation curves $c_\ell(r)$ for different values of the Paganin-filter parameter ℓ , with $\ell_0 < \ell_1 < \ell_2 < \ell_3$. Each curve $c_\ell(r)$ must always remain above its corresponding lower bound $B_\ell(r)$. (b) NPS curve, given by W_ℓ , as a function of ℓ . As ℓ increases, the correlation curve C_ℓ always remains above B_ℓ , monotonically tending toward 1, while the tail of W_ℓ monotonically decays toward zero.

$$\phi(\ell) := \int_{r_{\min}}^1 \frac{W_\ell(r)}{\|W_\ell\|_\infty} dr \leq \epsilon.$$

Then $\ell_{r_{\min}}$ provides a data-driven, noise-aware choice of Paganin-filter strength that enforces the prescribed suppression of high-frequency noise while remaining consistent with the split-image FRC bound.

The functional $\phi(\ell)$ in Proposition 1 measures the residual normalized noise power in the high-frequency band $[r_{\min}, 1]$. Throughout this work we set $r_{\min} = 0.5$ as a conservative default, and we select $\ell_{r_{\min}}$ as the smallest ℓ satisfying $\phi(\ell) \leq \epsilon$. For convenience, we restate this condition as

$$\phi(\ell) := \int_{r_{\min}}^1 \frac{W_\ell(r)}{\|W_\ell\|_\infty} dr \leq \epsilon. \quad (7)$$

The NPS curve given by W_ℓ , as a function of the Paganin-filter parameter ℓ , is shown in schematic form in Fig. 1(b). As the value of ℓ increases, the correlation curve C_ℓ always remains above B_ℓ [ensuring consistency, as expected (Miqueles *et al.*, 2025b)], monotonically tending toward unity while the tail of W_ℓ monotonically decays toward zero, so that the integral $\phi(\ell)$ given by equation (7) becomes progressively smaller.

We observe that the $\ell_{r_{\min}}$ such that $\phi(\ell) \leq \epsilon$ is obtained by a straightforward bisection method which, thanks to the monotonic behavior of $\phi(\ell)$ with respect to ℓ , typically converges in only a few iterations. As defined, the near-optimal value of ℓ derived from the NPS criterion depends explicitly on r_{\min} . In practical terms, the larger r_{\min} is, the lower the resulting ℓ must be; conversely, for smaller r_{\min} the solution tends toward higher values of ℓ . Hence, there is a conceptual trade-off between committing to a single value or to an *interval* for ℓ , reflecting our uncertainty about where the ‘high-frequency’ range should begin. We therefore opt to specify a *fuzzy* interval of near-optimal values, for example $\ell_{0.6} < \ell < \ell_{0.5}$, or, more generally,

$$\ell_{r_{\min}+\Delta r} \leq \ell \leq \ell_{r_{\min}}, \quad \Delta r > 0, \quad (8)$$

where the span Δr can be chosen in accordance with the spectral resolution (*i.e.* the ring partition) employed in computing $W_\ell(r)$ and $c_\ell(r)$. In principle, there may exist an additional criterion—defined via a suitable cost function—that selects a single ℓ^* within this interval.

2.2. Standard deviation and correlation criteria

To obtain a second criterion for evaluating the near-optimal parameter ℓ , we now consider two Paganin-filtered images p_ℓ and $p_{\ell+h}$, where h denotes a suitable step size in the filter parameter. It is natural to expect that, for sufficiently small variations of h , the difference between the images p_ℓ and $p_{\ell+h}$ remains small for certain values of ℓ . In this sense, as ℓ increases, noise is progressively more strongly suppressed and the difference $p_\ell - p_{\ell+h}$ mainly contains only numerical noise associated with the subtraction of the two images. On the other hand, for small values of ℓ , the same difference preserves structural details of the object. Thus, for a fixed step h , we define the criterion function

$$f(\ell) = \text{std} \left(\frac{\partial p_\ell}{\partial \ell} \right) \simeq \text{std}(p_\ell - p_{\ell+h}), \quad (9)$$

which measures the degree of noise suppression described above. In other words, the larger the standard deviation (std) between two images, the lower their mutual correlation. Based on this reasoning, we search for an optimal point of f in order to define a near-optimal parameter ℓ^* . In Appendix A we demonstrate that $f(\ell)$ is monotonically decreasing for large values of ℓ , but necessarily possesses a maximum in the interval $(0, \infty)$. To restrict the search space of this maximum, we focus on the interval $[0, \ell_{r_{\min}}]$, where $\ell_{r_{\min}}$ is the value obtained by the NPS-based criterion described in the previous section. Consequently, we may formulate the following optimization problem,

$$\ell_{\text{std}} = \arg \max \{ f(\ell) : \ell \in [0, \ell_{r_{\min}}] \}. \quad (10)$$

A schematic illustration of this function is presented in Fig. 2, where the elbow of f is clearly visible, corresponding to the near-optimal point ℓ_{std} .

Alternatively, as a counterpart to the standard deviation function, we may define another merit function via

$$g(\ell) = \left\| \frac{\partial p_\ell}{\partial \ell} \right\| \simeq \frac{\langle p_\ell, p_{\ell+h} \rangle}{\|p_\ell\| \|p_{\ell+h}\|}, \quad (11)$$

where (as previously mentioned) $\|\cdot\|$ denotes a norm induced by an inner product $\langle \cdot, \cdot \rangle$ in the function space L_2 . In Appendix B we provide a detailed derivation of the approximation in equation (11). Since g is a measure of the correlation between consecutive images, it is expected, in the same spirit as the function f , that g possesses a local minimum. Indeed, as $g \rightarrow 1$ when $\ell \rightarrow 0$ and $g \rightarrow 1$ as $\ell \rightarrow \infty$, it is natural to expect that g admits a local minimum in the interval $(0, \infty)$. A schematic plot of the function g is illustrated in Fig. 2, where the near-optimal value ℓ_{corr} is obtained through

$$\ell_{\text{corr}} = \arg \min \{ g(\ell) : \ell \in [0, \ell_{r_{\min}}] \}. \quad (12)$$




Figure 2 Typical shape of the curves $\{f, g\}$ in equation (9), exhibiting the optimal point within the interval $(0, \ell_{r_{\min}})$.




Figure 3 Representative projection frames illustrating the versatility of the proposed approach: (a) a natural rock sample, (b) a wooden toothpick with a micro gear inside, and (c) a thin biological tissue (soft murine cardiac tissue). Data were collected at the MOGNO beamline of the Sirius synchrotron light source.

3. Experiments

The numerical experiments consider the images shown in Fig. 3. The first image (a) represents a rock sample, whereas the second image (b) corresponds to a wood specimen in which a titanium object is embedded. The third image is a biological soft tissue (soft cardiac tissue from a murine sample). All datasets were acquired at an X-ray energy of 22 keV, with effective pixel size $\tau = 1.33 \mu\text{m}$ and sample-to-detector propagation distance of $z = 0.08 \text{ m}$ for sample (a); $\tau = 0.6 \mu\text{m}$ and $z = 0.4 \text{ m}$ for sample (b); and $\tau = 1.22 \mu\text{m}$, $z = 0.15 \text{ m}$ for sample (c). In Fig. 3, the histogram of each image is displayed in the lower-left corner of its respective panel, allowing one to visualize the separation (or overlap) between the intensity distributions. The images shown correspond to the normalization of the raw data, *i.e.* $-\log(r)$ with $r = I/I_0$. According to equation (1), these images are equivalent to applying the Paganin filter with $\ell = 0$. From a practical standpoint, the goal is to enhance contrast via the regularization operator in equation (1) while simultaneously minimizing the intrinsic measurement noise. This contrast gain typically facilitates subsequent data-analysis tasks such as segmentation and phase quantification, since image quality can be assessed, for example, by the separation of histogram peaks and by improved edge definition. Accordingly, ℓ is expected to be chosen so as to promote such separation without introducing an excessive loss of high-frequency detail. In what follows, we report and compare the near-optimal values obtained by the proposed criteria (based on NPS, standard deviation, and correlation), and discuss the visual and quantitative impact of this choice for the studied samples.

We begin by presenting the behavior indicated in Figs. 1(a) and 1(b), now using the real data. Fig. 4 shows the empirical behavior of the FRC curves $c(r)$ given by equation (5), computed via the same ring-wise radial partition in frequency space adopted in the theoretical analysis. In this figure we report, for reference, the value $\ell_{r_{\min}}$ obtained with $r_{\min} = 0.5$ (see Table 1), as well as the unfiltered case $\ell = 0$. This allows a side-by-side comparison of the three quantities of interest $\{c_\ell, W_\ell, B_\ell\}$: the FRC curve, the NPS curve, and the consistency lower bound. As predicted by theory (see the previous

section), as ℓ increases, the curve $c_\ell(r)$ rises and departs from $B_\ell(r)$ across the entire frequency range, satisfying the consistency condition $c_\ell(r) > B_\ell(r)$ (Miqueles *et al.*, 2025b). In parallel, the NPS $W_\ell(r)$ exhibits progressive suppression at high spatial frequencies, with its ‘tail’ becoming flatter and of smaller magnitude. For a fair comparison, the curves are obtained by averaging over each ring S_r and are displayed on the same discretization of $r \in [0, 1]$; when appropriate, W_ℓ is normalized by a suitable norm so that different scales are directly comparable across samples.

By requiring $W_\ell(r)$ to decay smoothly toward values close to zero from $r_{\min} = 0.5$ onward—that is, by enforcing the criterion $\phi(\ell) \leq \epsilon$ in equation (7)—we ensure that high-frequency noise contributions are effectively suppressed without unduly penalizing the lower-frequency components that contain the relevant structural information. In this way, the selected $\ell_{r_{\min}}$ simultaneously satisfies (i) $c_\ell > B_\ell$ (ensuring internal consistency) and (ii) high-frequency noise mitigation, thereby characterizing the near-optimality of the parameter under the proposed criterion. We further note that the choice of r_{\min} can be tuned according to the spectral resolution of the ring partitioning or the expected scale of meaningful contrast variations, defining a ‘fuzzy’ interval of plausible ℓ values that adapts naturally to the experimental conditions and detector sampling [*cf.* equation (8)]. This flexibility makes the procedure robust to moderate variations in signal-to-noise ratio and acquisition geometry.

We observed that the bisection method used to determine the near-optimal value $\ell_{r_{\min}}$ typically takes about 100 s for a search over the interval $[0, 10^8]$, which we consider acceptable given the large dynamic range of the parameter. By recursively halving the search range into smaller subintervals, the procedure consistently converges to an appropriate ℓ value within a few dozen iterations, without the need for manual tuning. In practice, the algorithm is highly stable and reproducible across datasets of different sizes and exposure levels, reinforcing the reliability of the near-optimal selection rule.

To obtain the values of ℓ_{std} and ℓ_{corr} , we construct the curves f and g defined in equations (9) and (11), respectively, over a logarithmically spaced set of ℓ values (see Fig. 5). In practice, we compute f, g on a mesh of 128 points split into eight

Table 1

Estimated values of the Paganin-filter parameter ℓ , obtained using the two methods proposed in this work for determining the near-optimal ℓ in equation (1).

Sample (a)	$\ell_{r_{\min}} = 381.84 \text{ m}^2$	$\ell_{\text{std}} = 62.65 \text{ m}^2$	$\ell_{\text{corr}} = 62.70 \text{ m}^2$
Sample (b)	$\ell_{r_{\min}} = 2402.81 \text{ m}^2$	$\ell_{\text{std}} = 304.39 \text{ m}^2$	$\ell_{\text{corr}} = 540.35 \text{ m}^2$
Sample (c)	$\ell_{r_{\min}} = 3203.74 \text{ m}^2$	$\ell_{\text{std}} = 42.18 \text{ m}^2$	$\ell_{\text{corr}} = 75.72 \text{ m}^2$

logarithmic subintervals, which provides both sufficient density and computational efficiency to reveal the maxima and minima of interest. This layout resolves the intricate peak–valley structure of f and g that a naive linear sampling would often miss or undersample—especially near very small or very large ℓ , where the system response changes rapidly (by several orders of magnitude). Building the curves f, g amounts to a single computational loop that depends only on the normalized frame $r = II_0$: for each ℓ , we apply the filtering operations




Figure 4
Curves $\{c_\ell, W_\ell, B_\ell\}$ corresponding to the FRC, the NPS, and the consistency lower bound, respectively, for each sample (a), (b) and (c). As ℓ increases, c_ℓ exceeds B_ℓ (consistency), while W_ℓ approaches zero at high spatial frequencies.

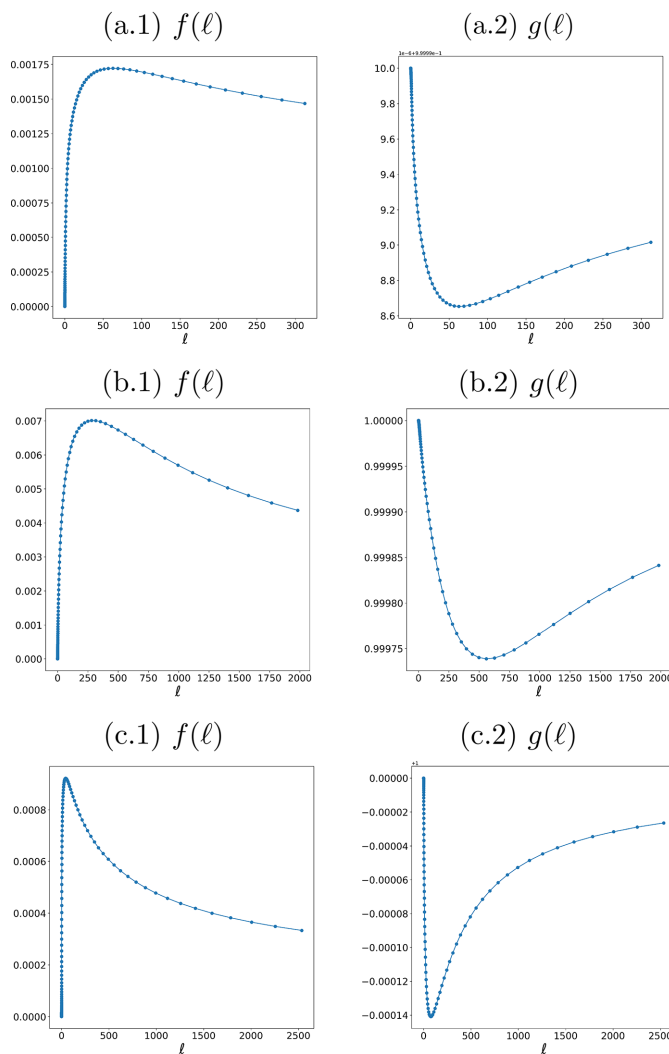


Figure 5
Curves $\{f, g\}$ defined in equations (9) and (11) [for samples (a), (b) and (c)], evaluated over log-spaced grids of $\ell \in [0, \ell_{r_{\min}}]$. Column (1) shows $f(\ell)$ and column (2) shows $g(\ell)$. The maximum of f and the minimum of g yield the near-optimal values ℓ_{std} and ℓ_{corr} , respectively.

to obtain p_ℓ and $p_{\ell+h}$, as specified in equations (9) and (11). For our datasets, this loop takes approximately 2 min on a standard CPU workstation, which we regard as a reasonable computational cost for selecting a near-optimal ℓ . Moreover, the computation is embarrassingly parallel (Breen *et al.*, 2008) and can be further accelerated on multi-core or GPU hardware if required for batch processing or high-throughput analysis.

For sample (a), the Paganin-filtered images obtained with the parameters ℓ_{std} and $\ell_{r_{\min}}$ [cf. equation (1)] are shown in Figs. 6(a.1) and 6(a.2), respectively. The central sectional line profiles are plotted beneath each panel, extracted along the midline of the image to enable a direct, pointwise comparison of edge sharpness and background fluctuations. For sample (b), the corresponding results are displayed in Figs. 6(b.1) and 6(b.2). Sample (c) is shown in Figs. 6(c.1) and 6(c.2). In every panel, the histogram of the displayed image appears in the lower-left corner and can be directly compared with the raw-

data histograms in Fig. 3. Consistent with the proposed selection criteria, the histograms exhibit increased separability between intensity modes after filtering, which facilitates downstream segmentation and thresholding after tomographic reconstruction—a key practical advantage of the Paganin-filter approach, particularly in automated workflows.

In our data, $\ell_{r_{\min}}$ yields stronger high-frequency suppression, while ℓ_{std} preserves slightly more fine detail, leading to a

balanced trade-off between denoising and edge retention. The overlaid line profiles further demonstrate a clear reduction in background variance and smoother plateau regions, with only modest attenuation of steep transitions. Overall, these results corroborate the proposition that the chosen parameter values significantly improve image contrast and uniformity while mitigating high-frequency noise, thereby enhancing interpretability and robustness for subsequent quantitative analysis and feature extraction.

4. Discussion

It is worth recalling that the physical meaning of the Paganin-filter parameter ℓ is to remove phase-contrast artifacts so as to recover a projected thickness of the specimen (Paganin *et al.*, 2002). In the single-distance, homogeneous TIE formulation, this corresponds to choosing ℓ to be large enough such that the phase term is effectively retrieved and the resulting image behaves as a projected thickness (up to a proportionality factor). In practice, however, such values of ℓ are typically very large and the associated low-pass kernel $\hat{K}_\ell(\mathbf{q}) = (1 + \ell\|\mathbf{q}\|^2)^{-1}$ becomes extremely aggressive at high spatial frequencies, producing very smooth images and, consequently, a noticeable loss of spatial resolution relative to the measured data. For instance, for sample (a) the thickness-recovery choice yields $\ell = 2793 \text{ m}^2$, whereas for sample (b) it yields $\ell = 56810 \text{ m}^2$. These numbers can also be determined numerically via a simple bisection search over ℓ for which the filtered, log-normalized image p_ℓ attains a positive (or nonnegative) signal profile along a representative line section, *i.e.* one increases ℓ while monitoring the filtered profile until a prescribed positivity criterion is met. While this procedure is physically meaningful for projected-thickness recovery, the resulting ℓ values are orders of magnitude larger than those in Table 1; by design they prioritize strong phase-contrast-fringe removal over resolution preservation and are therefore excluded from the comparative analysis in this study.

A second point concerns how the parameter values reported in Table 1 impact the tomographic reconstruction itself. It is well known that reconstruction algorithms—whether analytical (*e.g.* filtered backprojection) or iterative—introduce numerical artifacts such as streaks, ring artifacts, residual beam-hardening signatures, and edge over-/undershoots. These effects can distort the final histogram of the reconstructed slices or volumes, thereby hampering subsequent segmentation and quantitative analysis. Encouragingly, with the numerical selection strategy proposed here, the reconstructions preserve essentially a similar level of histogram separability achieved at the projection stage: the high-frequency noise is mitigated without excessive blurring, and the class boundaries in intensity space remain well separated after backprojection. This behavior is illustrated in Fig. 7, where we compare two reconstructions obtained from identical data: (i) without the Paganin filter and (ii) with the filter using a near-optimal ℓ . In addition to the visual impression, the improvement is corroborated by the tighter histograms and smoother midline profiles, which are consis-

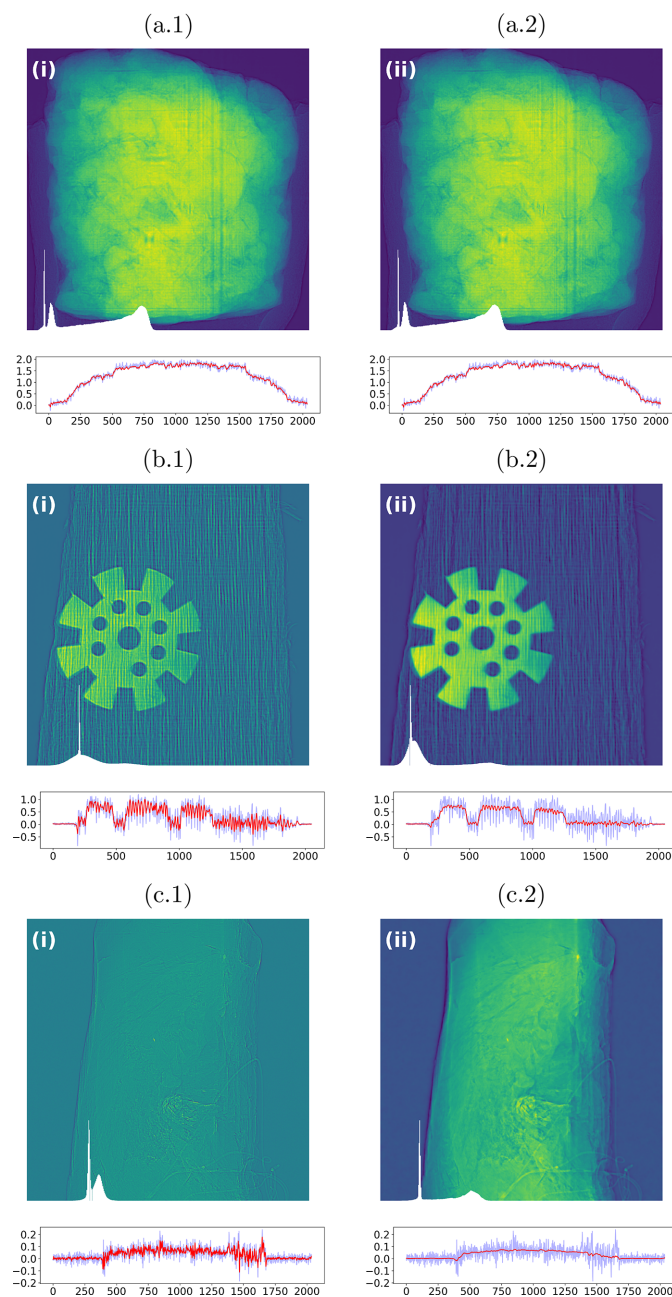


Figure 6 Paganin-filtered images corresponding to the parameter values ℓ_{std} and $\ell_{r_{\min}}$ listed in Table 1. For each image, the central line profile is shown beneath the panel, and the corresponding histogram appears in the lower-left corner, aiding visual assessment of contrast and noise. Accompanying line-profile plots show the noisy unfiltered profile (blue line) and the filtered profile obtained using the near-optimal proposed filter value (red line).




Figure 7 Tomographic reconstructions for sample (a), without (i) and with (ii) the Paganin filter using the near-optimal value ℓ_{std} , which improves histogram separation.

tent with our NPS- and correlation-based criteria. In short, the selected ℓ balances phase-artifact suppression and resolution retention so that the advantages observed at the projection level carry through to the reconstructed images, facilitating robust thresholding and phase segmentation downstream.

Another way to quantify the benefit of using the near-optimal values of ℓ is through computation of the standard deviation relative to the mean intensity for each of the filtered images. In practice, for a selected region of interest (ROI) within the reconstructed image, we evaluate a local deviation map that represents, over a coarse grid, the spatial distribution of intensity fluctuations. This map provides an intuitive visual indicator of the performance gain obtained when using each of the parameter values ℓ_{std} , $\ell_{r_{\text{min}}}$. The local deviations quantify how much the gray levels vary around their local mean, thus serving as a measure of local contrast (or texture) rather than image sharpness. In homogeneous areas—where intensity variations are small—the local deviation map exhibits low values, indicating smooth regions with little texture. Conversely, in regions containing edges, fine structural details, or strong noise components, the deviation values increase, yielding a higher contrast response.

These local contrast maps are displayed in Fig. 8. In this figure, column (i) corresponds to the reconstructed regions of interest obtained using $\ell = 0$, ℓ_{std} , and $\ell_{r_{\text{min}}}$ for sample (b). A gradual increase in image blur can be visually perceived from above to below. In column (ii), we present the corresponding local contrast maps derived from column (i), respectively. The color scale of these maps highlights the degree of local intensity dispersion, where brighter tones indicate larger fluctuations. As observed, the case $\ell = 0$ produces an image with lower contrast and weaker local variations, whereas $\ell_{r_{\text{min}}}$ results in stronger local contrast at the expense of increased blurring due to the higher suppression of high frequencies. Between these two extremes, ℓ_{std} offers the most balanced compromise, effectively enhancing local contrast while maintaining acceptable sharpness—thereby facilitating quantita-

tive image analysis, segmentation, and texture-based classification in subsequent processing stages.

The overall trend of Fig. 8 remains consistent: the near-optimal parameter $\ell_{r_{\text{min}}}$ increases the visual separability of features and enhances the apparent texture contrast, while ℓ_{std} provides an intermediate response that is particularly suitable for quantitative tomographic reconstruction workflows, where both noise control and structural fidelity are required. Together, these observations reinforce the robustness and general applicability of the proposed selection criterion across samples of different physical nature and contrast regime.

A complementary view is provided by the local contrast analysis summarized in Fig. 8. Because the maps are based on the standard deviation relative to the local mean (a blockwise coefficient of variation), they respond primarily to textural fluctuations rather than to high-frequency edge sharpness. This distinction is useful in practice: it decouples *perceived* contrast improvement from resolution loss, making explicit the trade-off negotiated by the Paganin-filter parameter. In our datasets, $\ell_{r_{\text{min}}}$ reliably boosts local contrast in heterogeneous regions (*e.g.* porous or fibrous textures), while ℓ_{std} prevents over-smoothing of steep transitions, preserving diagnostically relevant edges. As a working guideline, we therefore recommend choosing the parameter ℓ_{std} when edge delineation is critical (*e.g.* metrology or thin-feature quantification) and $\ell_{r_{\text{min}}}$ when robust texture separation or histogram mode separability is the priority (*e.g.* automated segmentation pipelines).

From an implementation standpoint, both selection procedures are reproducible and computationally lightweight. The NPS-based bisection is monotone in $\phi(\ell)$ and needs only ring-wise frequency statistics; the $\{f, g\}$ scans are embarrassingly parallel over a log-spaced grid and depend solely on applying p_ℓ and $p_{\ell+h}$. In practice, this design makes the selection step easy to cache and to amortize across batches of projections or slices. We note three limitations that suggest avenues for future refinement: (i) non-stationary noise and detector




Figure 8

Local-contrast assessment for near-optimal Paganin-filter parameters using sample (b) as a reference. Column (i) shows ROIs reconstructed with $\ell = 0$, ℓ_{std} , and $\ell_{r_{\text{min}}}$, respectively; column (ii) gives corresponding blockwise local-contrast maps (standard deviation relative to the local mean), where brighter tones indicate larger intensity dispersion. Consistently, $\ell = 0$ gives lower contrast, $\ell_{r_{\text{min}}}$ yields higher contrast with stronger blur, and ℓ_{std} provides the best compromise for downstream quantitative analysis.

anisotropy may violate the ‘isotropic ring’ assumption underlying $W_{\ell}(r)$; (ii) strong ring or beam-hardening artifacts can bias both the NPS tail and the local-contrast estimates unless pre-corrected; and (iii) the effective modulation transfer function (MTF) of the acquisition system is not explicitly modeled here. Incorporating an MTF-aware normalization or adaptive (angle-dependent) ring partitioning could further sharpen the link between $\phi(\ell)$ and visually perceived resolution. Despite these caveats, the observed agreement between the NPS-consistency criterion and the $\{f, g\}$ extrema across diverse specimens (rock, plant, and soft tissue) supports the robustness and generality of the approach.

The near-optimal Paganin-filter parameter introduced in this work should not be interpreted as relating to an intrinsic physical property of the sample. Rather, it provides an effective,

data-driven estimate of the regularization strength required by the Paganin filter when prior knowledge of the material parameters is incomplete or uncertain. In the homogeneous-material approximation underlying the Paganin method, the filter parameter is directly related to the ratio β/δ , which itself encodes assumptions regarding composition, density, and beam energy. In practice, however, these quantities are often only approximately known, and small deviations can significantly affect the balance between noise suppression and spatial resolution. Important context for the preceding statement is provided by the well known tradeoff between noise and spatial resolution, both in general imaging-physics terms (den Dekker & van den Bos, 1997; Neifeld, 1998; van den Bos & den Dekker, 2001) and in the specific setting of inverse problems in X-ray propagation-based phase-contrast imaging (Gureyev *et al.*, 2014b; de Hoog *et al.*, 2014; Gureyev *et al.*, 2015; Gureyev *et al.*, 2016; Gureyev *et al.*, 2020). The proposed criterion can be viewed as compensating for the previously mentioned uncertainties by selecting a reconstruction parameter that stabilizes the phase-retrieval process while remaining consistent with the physical constraints of the model. Although the resulting parameter may be interpreted in physical terms of an effective strength associated with the degree of propagation-based X-ray phase contrast (Snigirev *et al.*, 1995; Cloetens *et al.*, 1996; Wilkins *et al.*, 1996), its primary role is computational, *i.e.* to adapt the regularization to the data in a physically meaningful way, without introducing additional free parameters or significant computational overhead.

The sensitivity of the proposed parameter-selection criterion to experimental conditions such as noise level, exposure time, and detector binning follows naturally from its formulation. Changes in exposure time or photon flux directly affect the signal-to-noise ratio of the acquired projections, which is captured by the FRC computed between independently split images. In photon-starved regimes, increased noise leads to a faster decay of the FRC at high spatial frequencies, resulting in the selection of a more conservative regularization parameter that prioritizes noise suppression. Conversely, in high-flux or long-exposure conditions, improved photon statistics allow higher spatial frequencies to be preserved, yielding a weaker regularization. Detector binning produces a similar effect by modifying the effective noise level and frequency content of the data. As a result, while the proposed criterion does not explicitly model these acquisition parameters, it adapts implicitly to different experimental regimes through the data-driven FRC constraint, remaining applicable across a wide range of synchrotron imaging conditions.

Although the proposed criterion is formulated within the homogeneous-material assumption underlying the original form of the Paganin method, its data-driven nature allows it to remain meaningful in more general imaging scenarios. In objects with spatially varying β/δ ratios or heterogeneous material composition, the selected parameter can be interpreted as an effective regularization that balances noise suppression and resolution across the field of view rather than as a locally exact physical quantity. In this sense, the method

does not resolve material heterogeneity, but provides a stable global parameter choice that is robust to moderate deviations from homogeneity. In multi-distance phase-retrieval setups, the criterion may be applied independently at each propagation distance, or used to inform a consistent parameter selection when combining information from multiple distances. While dedicated multi-material or multi-distance methods remain necessary for quantitative material discrimination, the proposed approach offers a complementary and computationally inexpensive tool for parameter selection in practical experimental conditions.

5. Conclusion

We have introduced two quantitative and complementary criteria to select a near-optimal Paganin-filter parameter for propagation-based phase-contrast X-ray tomography, replacing empirical tuning with a reproducible and noise-aware procedure. The first criterion relies on the noise power spectrum evaluated over FRC-style radial bands and enforces both internal consistency—via the lower bound $B_\ell(r)$ with $c_\ell(r) > B_\ell(r)$ —and high-frequency noise suppression through the integral constraint $\phi(\ell) \leq \epsilon$. The second criterion exploits the extremal behavior of two merit functions derived from finite-difference approximations to $\partial p_\ell / \partial \ell$: the standard-deviation curve $f(\ell)$ (maximized at ℓ_{std}) and the correlation curve $g(\ell)$ (minimized at ℓ_{corr}). Across rock, plant, and soft-tissue datasets, the three estimates ($\ell_{r_{\text{min}}}$, ℓ_{std} , ℓ_{corr}) produced reconstructions with improved histogram separability, lower background variance, and preserved structural interpretability.

From a practical standpoint, $\ell_{r_{\text{min}}}$ reliably achieves stronger high-frequency suppression, which benefits segmentation and texture separation, while ℓ_{std} typically preserves slightly sharper edges and fine detail—often the best compromise when downstream metrology or boundary localization is critical. The local-contrast maps corroborate these observations by decoupling textural gain from edge sharpness, clarifying the trade-off negotiated by ℓ . Computationally, both selection paths are lightweight: the NPS-based bisection converges in a few iterations, and the (f, g) scans can be evaluated on a log-spaced grid in minutes on a CPU, with straightforward opportunities for parallel/GPU acceleration.

The proposed methodology is simple to implement and robust to moderate variations in both signal-to-noise ratio and acquisition geometry. It integrates naturally with standard reconstruction pipelines, carrying the projection-level benefits into the volume domain without excessive blurring or artifact amplification. Limitations include the implicit assumption of isotropic, stationary noise and the lack of an explicit MTF model; both can be addressed in future work by incorporating angle-dependent ring partitioning, MTF-aware normalization, or adaptive r_{min} selection. Finally, the accompanying open-source implementation codes (see Appendix C) provide a transparent path to adoption and benchmarking, supporting routine, objective parameter selection in synchrotron-based tomographic workflows.

APPENDIX A

Optimal existence

Let $f(\ell)$ be defined in equation (9) and consider $p_\ell = -\log P_\ell(r)$ with $P_\ell = (I - \ell\Delta)^{-1}$. By defining $u_\ell := P_\ell(r)$ we can compute the derivative of p_ℓ with respect to ℓ through

$$\frac{\partial p_\ell}{\partial \ell} = -\frac{1}{u_\ell} \frac{\partial u_\ell}{\partial \ell} = -\frac{(I - \ell\Delta)^{-1} \Delta u_\ell}{u_\ell}. \quad (13)$$

For linear steps $h > 0$, we define

$$f_h(\ell) = \text{std}(p_\ell - p_{\ell+h}) \quad (14)$$

so that by Taylor expansion we obtain

$$p_{\ell+h} - p_\ell = h \left. \frac{\partial p_\ell}{\partial \ell} \right|_{\ell} + \mathcal{O}(h^2), \quad \left. \frac{\partial p_\ell}{\partial \ell} \right|_{\ell=0} = -\frac{\Delta r}{r}. \quad (15)$$

It follows therefore that

$$f_h(0) = \text{std}(p_0 - p_h) = |h| \text{std}\left(\frac{\Delta r}{r}\right) + \mathcal{O}(h^2), \quad (16)$$

that is, for fixed and nonzero h , the value of $f_h(0)$ is of order h . In the limit $h \rightarrow 0$, we have $f_h \rightarrow 0$. On the other hand, from equation (13), we know that $\|(I - \ell\Delta)^{-1} \Delta\| \leq 1/\ell$, so that

$$p_{\ell+h} - p_\ell = \int_\ell^{\ell+h} \frac{\partial p_\tau}{\partial \tau} d\tau \quad (17)$$

implies

$$\|p_{\ell+h} - p_\ell\|_2 \leq \int_\ell^{\ell+h} \frac{C}{\tau} d\tau = C \log(1 + h/\ell) \quad (18)$$

for some bounding constant C . It then follows that, when $\ell \rightarrow \infty$, we can say

$$f_h(\ell) = \mathcal{O}\left[\log\left(1 + \frac{h}{\ell}\right)\right] = \mathcal{O}\left(\frac{1}{\ell}\right).$$

Since f_h is continuous in $\ell \geq 0$ (as it involves a set of continuous operations), and reaches null values at the extremes of the interval $[0, \infty)$, it follows that f_h necessarily attains a local maximum in this interval. Moreover, as $f \simeq f_h$ for small h , the same behavior of f_h can be transferred to the function f , so that it also attains a local maximum in the interval $[0, \infty)$.

APPENDIX B

Cosine similarity

Let $\phi(\ell) = p_\ell / \|p_\ell\|_2$. To define the approximation in equation (11), it is enough to note that

$$\frac{\langle p_\ell, p_{\ell+h} \rangle}{\|p_\ell\|_2 \|p_{\ell+h}\|_2} = \langle \phi(\ell), \phi(\ell+h) \rangle. \quad (19)$$

Thus, by applying the Taylor expansion of ϕ in ℓ , and using the fact that $\langle \phi, \phi \rangle = 1$ and $\langle \phi, \phi' \rangle = 0$, we obtain

$$g(\ell) = \langle \phi(\ell), \phi(\ell+h) \rangle = 1 - \frac{h^2}{2} \|\phi'(\ell)\|_2^2 + \mathcal{O}(h^3),$$

from which the proposed approximation follows. It is worth noting that this approximation is also referred to as ‘cosine similarity’, for obvious reasons.

APPENDIX C

Python codes

The codes implementing the theoretical framework presented in this manuscript are included here for completeness and to ensure full scientific reproducibility. They are concise and self-contained, being designed so as not to interrupt the flow or clarity of the main text. The Python function `<optimal_nps>`, shown in Code 1, implements the computation of the near-optimal value $\ell_{r_{\min}}$ —as defined by the criterion in equation (7)—for a given input parameter r_{\min} , together with the normalized image $r = I/I_0$, where I_0 denotes the measured flat-field and I denotes the corresponding transmission recorded in the experiment. For the datasets analyzed in this work, r corresponds to a real-valued image of 2048×2048 pixels. On a standard CPU workstation, the calculation of $\ell_{r_{\min}}$ takes approximately 40 s on average, demonstrating that the procedure is computationally lightweight and practical for routine tomographic data processing.

Code 1 Computation for the optimal value $\ell_{r_{\min}}$

```
def pad(I, pad):
    return np.pad(I, pad_width=(pad, pad), (pad, pad)), mode='edge')

def unpad(I_padded, pad):
    if pad == 0:
        return I_padded
    return I_padded[pad:-pad, pad:-pad]

def _paganin_(r, L, log=True):
    _PAD_ = 100
    r = pad(r, _PAD_)
    Ny, Nx = r.shape
    ky = np.fft.fftfreq(Ny, d=1)
    kx = np.fft.fftfreq(Nx, d=1)
    k2 = np.add.outer(ky**2, kx**2)

    kernel = 1 / (1 + L * k2)
    r_hat = fft2(r)
    filtered_hat = kernel * r_hat
    u = ifft2(filtered_hat)

    p = -np.log(u)
    p = unpad(p, _PAD_)
    return p.real

def curves(r, L):
    ringwidth = 1
    P = _paganin_(r, L)
    u = P[::2, ::2]
    v = P[1::2, 1::2]

    def inprod(x, y):
        return (x * numpy.conjugate(y)).sum()

    U = numpy.fft.fftfreq( numpy.fft.fft2(u) )
    V = numpy.fft.fftfreq( numpy.fft.fft2(v) )

    dx = 2.0/(u.shape[1]-1)
    wc = 1.0/(2*dx)

    Area = (dx ** 2) * u.size

    t = numpy.linspace(-wc,wc,u.shape[0])
    wx,wy = numpy.meshgrid(t,t)
    wy = numpy.flipud(wy)

    r = numpy.linspace(0, wc, u.shape[1]/2)

    rc = ringwidth * (t[1]-t[0])

    curve = numpy.ones(r.shape) + 0*1j
    T = numpy.ones(r.shape)
    Ts = numpy.ones(r.shape)
    B = numpy.ones(r.shape)
    NPS = np.zeros(r.shape)

    for k in range(1,len(r)):

        mask = (wx**2 + wy**2 <= (r[k]+rc)**2) & (wx**2 + wy**2 >= (r[k]-rc)**2)
        Uat = U[mask]
        Vvat = V[mask]

        nUat = numpy.linalg.norm(Uat)
        nVvat = numpy.linalg.norm(Vvat)

        curve[k] = inprod(Uat, Vvat) / (nUat * nVvat)

        npixels = (mask.flatten()*1.0 > 0).sum()
        T[k] = ( 0.2071 + 1.9102/numpy.sqrt(npixels) )
        T[k] = T[k] / (1.2071 + 0.9102/numpy.sqrt(npixels))
        Ts[k] = 3.0 / numpy.sqrt(npixels/2.0)
```

```
M1 = numpy.max( abs( Uat ) )
M2 = numpy.max( abs( Vvat ) )
M = max(M1, M2)
B[k] = 2.0 * numpy.sqrt(M) / (1 + M)
Diff = Uat - Vvat
NPS[k] = np.sum(np.abs(Diff)**2) / (Area * npixels)

B[0] = 0
NPS[0] = 0

_curves_ = {}
_curves_['r'] = r/r.max()
_curves_['G'] = curve.real
_curves_['H'] = T
_curves_['S'] = Ts
_curves_['B'] = B
_curves_['NPS'] = NPS / NPS.max()
_curves_['diff'] = B[-1] - curve[-1].real

return _curves_

def optimal_nps( ratio, RMIN=0.5 ):
    def NPSCurve(u, v, ringwidth):
        #Assume squared images, with the same size
        def inprod(x, y):
            return (x * numpy.conjugate(y)).sum()

        U = numpy.fft.fftfreq( numpy.fft.fft2(u) )
        V = numpy.fft.fftfreq( numpy.fft.fft2(v) )
        dx = 2.0/(u.shape[1]-1)
        wc = 1.0/(2*dx)
        Area = (dx ** 2) * u.size
        t = numpy.linspace(-wc,wc,u.shape[0])
        wx,wy = numpy.meshgrid(t,t)
        wy = numpy.flipud(wy)
        r = numpy.linspace(0, wc, u.shape[1]/2)
        rc = ringwidth * (t[1]-t[0])
        NPS = np.zeros(r.shape)

        for k in range(1,len(r)):

            mask = (wx**2 + wy**2 <= (r[k]+rc)**2) & (wx**2 + wy**2 >= (r[k]-rc)**2)
            Uat = U[mask]
            Vvat = V[mask]

            nUat = numpy.linalg.norm(Uat)
            nVvat = numpy.linalg.norm(Vvat)
            npixels = (mask.flatten()*1.0 > 0).sum()

            # NPS
            Diff = Uat - Vvat
            NPS[k] = np.sum(np.abs(Diff)**2) / (Area * npixels)

        NPS[0] = 0

        curves = {}
        curves['r'] = r/r.max()
        curves['NPS'] = NPS / NPS.max()

    return curves

def eval_lambda_per_NPS(dic, eps=1e-4, rmin=0.5, rmax=1.0):
    r = np.asarray(dic['r'])
    nps = np.asarray(dic['NPS'])
    mask = (r >= rmin) & (r <= rmax)
    r_sub = r[mask]
    nps_sub = nps[mask]
    integral = np.trapz(nps_sub, r_sub)
    return integral - eps

def bisection_NPS(r, rmin, lambda_min=0, lambda_max=1000, eps=1e-4, tol=1e-6,
max_iter=30):
    ringwidth = 1

    for _ in range(max_iter):
        lambda_mid = 0.5 * (lambda_min + lambda_max)

        P = _paganin_(r, lambda_mid)
        u = P[::2, ::2]
        v = P[1::2, 1::2]
        dic = NPSCurve(u, v, ringwidth)

        val = eval_lambda_per_NPS(dic, eps, rmin=rmin)

        if abs(val) < tol:
            return lambda_mid

        if val > 0:
            lambda_min = lambda_mid
        else:
            lambda_max = lambda_mid

        if verbose:
            print(lambda_min, lambda_mid, lambda_max)

    return lambda_mid

#
return bisection_NPS( ratio, RMIN, lambda_min=0, lambda_max=1e+10, eps=1e-4,
tol=1e-6, max_iter=30)
```

The Python function `<optimal_std_corr>`, presented in Code 2, implements the computation of the near-optimal value of ℓ_{std} , ℓ_{corr} according to equations (9) and (11). This implementation follows directly from the theoretical formulation introduced in Section 2, providing a practical tool to determine ℓ through the joint evaluation of the standard-deviation and correlation criteria. The procedure scans the parameter range defined by $[0, \ell_{r_{\min}}]$ and identifies the

maximum and minimum of the merit functions $f(\ell)$ and $g(\ell)$, respectively, thereby yielding ℓ_{std} and ℓ_{corr} . The average execution time of the presented implementation is approximately 89 s when run solely on a CPU, without any parallelization or hardware acceleration. Despite this modest computational cost, the method remains efficient and reproducible, and can be further accelerated through straightforward parallelization strategies.

Code 2 Computation for the optimal values $\{\ell_{\text{std}}, \ell_{\text{corr}}\}$

```
def optimal_std_corr( ratio, lnps ):
    def meshLogspace(a, b, N, M):
        n_per_block = [N // M] * M
        for i in range(N % M):
            n_per_block[i] += 1
        edges = np.logspace(np.log10(a), np.log10(b), M+1)
        mesh = []
        for i in range(M):
            npts = n_per_block[i]
            block = np.logspace(np.log10(edges[i]), np.log10(edges[i+1]), npts,
                               endpoint=False)
            mesh.append(block)
        mesh.append((b))
        return np.concatenate(mesh)
    nL = 128
    nS = 8
    Ls = meshLogspace(0.001, lnps, nL, nS)
    f = numpy.zeros(nL-1)
    g = numpy.zeros(nL-1)
    for k in range(nL-1):
        pag = nps._paganin( ratio, Ls[k] )
        pagNext = nps._paganin( ratio, Ls[k+1] )
        f[k] = numpy.std( (pagNext - pag) )
        g[k] = (pagNext * pag).sum() / ( (pagNext**2).sum()**0.5 *
                                         (pag**2).sum()**0.5 )
    idxf = numpy.argmax(f)
    idxg = numpy.argmin(g)
    lopt_std = Ls [ idxf ]
    lopt_corr = Ls [ idxg ]
    return lopt_std, lopt_corr
```

Acknowledgements

This work was supported in part by the São Paulo Research Foundation (FAPESP) under Grant 2024/08590-6.

Funding information

The following funding is acknowledged: Fundação de Amparo à Pesquisa do Estado de São Paulo (grant No. 2024/08590-6).

References

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**, 1945–1948.
 Banterle, N., Bui, K. H., Lemke, E. A. & Beck, M. (2013). *J. Struct. Biol.* **183**, 363–367.
 Barrett, H. H. & Myers, K. J. (2004). *Foundations of Image Science*. Hoboken: John Wiley & Sons.
 Beltran, M. A., Paganin, D. M. & Pelliccia, D. (2018). *J. Opt.* **20**, 055605.
 Breen, B. J., Weidert, C. E., Lindner, J. F., Walker, L. M., Kelly, K. & Heidtmann, E. (2008). *Am. J. Phys.* **76**, 347–352.
 Cloetens, P., Barrett, R., Baruchel, J., Guigay, J.-P. & Schlenker, M. (1996). *J. Phys. D Appl. Phys.* **29**, 133–146.

Culley, S., Albrecht, D., Jacobs, C., Pereira, P. M., Leterrier, C., Mercer, J. & Henriques, R. (2018). *Nat. Methods* **15**, 263–266.
 Dainty, J. C. & Shaw, R. (1974). *Image Science: Principles, Analysis and Evaluation of Photographic-Type Imaging Processes*. London: Academic Press.
 de Hoog, F., Schmalz, G. & Gureyev, T. E. (2014). *Appl. Math. Lett.* **38**, 84–86.
 den Dekker, A. J. & van den Bos, A. (1997). *J. Opt. Soc. Am. A* **14**, 547–557.
 Eastwood, S. A., Paganin, D. M. & Liu, A. C. Y. (2011). *Opt. Lett.* **36**, 1878–1880.
 Goodman, J. W. (2000). *Statistical Optics*. New York: Wiley.
 Gureyev, T. E., Stevenson, A., Nesterets, Y. I. & Wilkins, S. (2004). *Opt. Commun.* **240**, 81–88.
 Gureyev, T. E., De Hoog, F., Nesterets, Ya. I. & Paganin, D. M. (2015). *ANZIAM J.* **55**, 1.
 Gureyev, T. E., Kozlov, A., Paganin, D. M., Nesterets, Ya. I. & Quiney, H. M. (2020). *Sci. Rep.* **10**, 7890.
 Gureyev, T. E., Mayo, S. C., Nesterets, Ya. I., Mohammadi, S., Lockie, D., Menk, R. H., Arfelli, F., Pavlov, K. M., Kitchen, M. J., Zanconati, F., Dullin, C. & Tromba, G. (2014a). *J. Phys. D Appl. Phys.* **47**, 365401.
 Gureyev, T. E., Nesterets, Y. I. & de Hoog, F. (2016). *Opt. Express* **24**, 17168–17182.
 Gureyev, T. E., Nesterets, Y. I., de Hoog, F., Schmalz, G., Mayo, S. C., Mohammadi, S. & Tromba, G. (2014b). *Opt. Express* **22**, 9087–9094.
 Gureyev, T. E., Nesterets, Ya. I., Kozlov, A., Paganin, D. M. & Quiney, H. M. (2017). *J. Opt. Soc. Am. A* **34**, 2251–2260.
 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.
 Miqueles, E. X., Ferraz, P. C., Lovric, G. & Paganin, D. M. (2025a). *IEEE Access* **13**, 128060–128070.
 Miqueles, E. X., Tonin, Y. R. & Luke, R. D. (2025b). *IEEE Trans. Comput. Imaging* **11**, 1047–1058.
 Neifeld, M. A. (1998). *Opt. Lett.* **23**, 1477–1479.
 Nesterets, Y. I. & Gureyev, T. E. (2014). *J. Phys. D Appl. Phys.* **47**, 105402.
 Nieuwenhuizen, R. P., Lidke, K. A., Bates, M., Puig, D. L., Grünwald, D., Stallinga, S. & Rieger, B. (2013). *Nat. Methods* **10**, 557–562.
 Paganin, D., Gureyev, T. E., Mayo, S. C., Stevenson, A. W., Nesterets, Y. I. & Wilkins, S. W. (2004). *J. Microsc.* **214**, 315–327.
 Paganin, D., Mayo, S. C., Gureyev, T. E., Miller, P. R. & Wilkins, S. W. (2002). *J. Microsc.* **206**, 33–40.
 Paganin, D. M., Favre-Nicolin, V., Mirone, A., Rack, A., Villanova, J., Olbinado, M. P., Fernandez, V., da Silva, J. C. & Pelliccia, D. (2020). *J. Opt.* **22**, 115607.
 Rieger, B., Droste, I., Gerritsma, F., ten Brink, T. & Stallinga, S. (2024). *Opt. Express* **32**, 21767–21782.
 Saxton, W. O. & Baumeister, W. (1982). *J. Microsc.* **127**, 127–138.
 Snigirev, A., Snigireva, I., Kohn, V., Kuznetsov, S. & Schelokov, I. (1995). *Rev. Sci. Instrum.* **66**, 5486–5492.
 Teague, M. R. (1983). *J. Opt. Soc. Am.* **73**, 1434–1441.
 Tortarolo, G., Castello, M., Diaspro, A., Koho, S. & Vicidomini, G. (2018). *Optica* **5**, 32–35.
 van den Bos, A. & den Dekker, A. (2001). *Adv. Imaging Electron. Phys.* **117**, 241–360.
 Wilkins, S. W., Gureyev, T. E., Gao, D., Pogany, A. & Stevenson, A. W. (1996). *Nature* **384**, 335–338.