



High-speed processing of X-ray wavefront marking data with the Unified Modulated Pattern Analysis (UMPA) model

FABIO DE MARCO,^{1,2,*}  SARA SAVATOVIĆ,^{1,2}  RONAN SMITH,³ 
VITTORIO DI TRAPANI,^{1,2}  MARCO MARGINI,^{1,2}  GINEVRA
LAUTIZI,^{1,2}  AND PIERRE THIBAUT^{1,2} 

¹*Department of Physics, University of Trieste, Via Valerio 2, 34127 Trieste, Italy*

²*Elettra-Sincrotrone Trieste, Strada Statale 14 – km 163.5, 34149 Basovizza, Italy*

³*Department of Physics, University of Southampton, University Road, Southampton SO17 1BJ, UK*

*fabiodomenico.demarco@units.it

Abstract: Wavefront-marking X-ray imaging techniques use e.g., sandpaper or a grating to generate intensity fluctuations, and analyze their distortion by the sample in order to retrieve attenuation, phase-contrast, and dark-field information. Phase contrast yields an improved visibility of soft-tissue specimens, while dark-field reveals small-angle scatter from sub-resolution structures. Both have found many biomedical and engineering applications. The previously developed Unified Modulated Pattern Analysis (UMPA) model extracts these modalities from wavefront-marking data. We here present a new UMPA implementation, capable of rapidly processing large datasets and featuring capabilities to greatly extend the field of view. We also discuss possible artifacts and additional new features.

© 2022 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

In conventional X-ray imaging, information retrieved from a sample is encoded purely by its attenuation of the incident light. X-ray phase-contrast imaging, on the other hand, retrieves information of an imaged sample by examining its effect on the phase of an incident wavefront. Since the attenuation of very thin samples, or samples composed of light elements, is often weak, resulting attenuation images can exhibit very low contrast. In such cases, X-ray phase-contrast techniques often provide images with a superior signal-to-noise ratio, since the real decrement δ of the index of refraction exceeds its imaginary part β by several orders of magnitude [1]. As X-ray detectors cannot determine the phase of a wavefront, X-ray phase-contrast imaging methods encode this information in an intensity pattern.

Over the last three decades, a wide range of methods for achieving such an encoding have been developed. Among these are “propagation-based imaging”, which exploits self-interference of the beam after being disturbed by a sample [2–5], “analyzer-based imaging”, which isolates wave components with a specific propagation direction through the use of an analyzer crystal [6,7], “grating-based imaging”, where the distortion of Talbot self-images produced by a transmission grating is analyzed [8–10], the closely-related, but non-interferometric “edge-illumination” technique [11–13], and finally a range of “wavefront-marking” methods, where a single optical element generates a directly resolved intensity modulation, whose distortion is then algorithmically analyzed. Besides attenuation and phase information, several of these methods [5,7,10,13] also retrieve “dark-field” information, which characterizes the amount of small-angle scattering.

A wide range of optical elements may be used for wavefront marking: besides ordered, periodic objects such as gratings [14,15], randomly-oriented structures such as sandpaper may also be used [16–18]. The latter approach produces (near-field) X-ray speckle [19], hence the method is also called “speckle-based imaging”. An important benefit of the method is its experimental

simplicity, and its ability to retrieve both differential-phase and dark-field information in two dimensions. Despite its relatively recent introduction, it has found applications in wavefront sensing, at-wavelength metrology of X-ray optics components, as well as biomedical and materials science imaging problems [20,21].

We here give a brief overview of existing approaches for extracting information from X-ray speckle imaging data. Two novel approaches have recently been demonstrated for extracting information from speckle-based X-ray data: firstly, the “ptychographic X-ray speckle tracking” (PXST) technique which iteratively reconstructs sample and wavefront from a series of images where the sample is laterally translated [22,23]. Secondly, from the insight that the Fokker-Planck differential equation can be used to model refraction and dark-field [24,25], “multi-modal intrinsic speckle-tracking” (MIST) has been developed. The distortion of speckle patterns is quantified by solving this equation [26], even allowing the extraction of directional scatter [27].

Earlier processing approaches instead attempt to correlate subregions of sample and reference images. Two reviews give a more in-depth insight to these different approaches [20,28]. They are derived from the differential interference contrast (DIC) technique, which uses laser speckle to characterize strain in a material [29].

In “X-ray speckle tracking” (XST), one image each is taken with and without the sample. To identify the displacement in a pixel of the sample image, a neighborhood of this pixel, the “(analysis) window”, is identified and the best match of this window in the reference image is sought by maximizing the zero-normalized cross-correlation (ZNCC) between the windows. The 2D displacement vector between these windows represents the local amount of refraction. Since the analysis window must be greater than one pixel, its extent limits the spatial resolution of resulting images [16,18].

In “X-ray speckle scanning” (XSS), the diffuser is displaced on a very fine 1D or 2D grid, yielding a 1D or 2D map of intensities for each pixel (with and without the sample). The displacement is calculated by identifying the shift between the two intensity maps. A higher spatial resolution and better angular phase sensitivity can be achieved, albeit at the expense of a slower acquisition and the need for very precise diffuser positioning [17].

In “X-ray speckle vector tracking” (XSVT), the diffuser is displaced in intervals of arbitrary shape, which yields a vector of intensities for each detector pixel (one value per diffuser position). For each sample image vector, a reference image vector is sought that maximizes the ZNCC. The shift is then given as the distance between the vectors’ pixel locations. Attenuation and dark-field data are also retrieved [30]. The method was later extended to also use an analysis window [31].

An alternative to maximizing the ZNCC is minimizing a cost function which, in addition to lateral shift, also incorporates attenuation and dark-field information. This approach was first applied in [32], and later generalized to the “Unified Modulated Pattern Analysis” (UMPA) model by the inclusion of diffuser stepping with arbitrary trajectories (like XSVT) [33].

An important limitation of most of the above-mentioned methods is the significant required computational effort. This is especially severe in the case of tomography, which requires a large number of projections. The problem is compounded by the steady increase in detector resolution and frame rate, and thus the generated data volume, especially at synchrotron beamlines.

In order to make the analysis of speckle imaging datasets more accessible, we here present an improved implementation of the UMPA model. Compared to the previous version [33], a number of new features have been introduced, and single-thread execution speed has been increased by two orders of magnitude.

Most importantly, the aspect of displacing the diffuser between measurements has been generalized to instead allow displacing the sample (similar to the technique presented in [22]). By defining a suitable sample “trajectory”, the effective field of view can be arbitrarily increased, limited only by the motor travel range. We demonstrate the package’s capabilities on sample-stepping measurements performed at two different synchrotron beamlines.

The speed increase is achieved by translation of the core routines to C++, while the ease of use of previous Python versions is maintained by embedding the C++ code in a Python interface through Cython [34]. Speed is further increased by parallelization with OpenMP.

We also discuss the presence of a bias in the estimation of retrieved speckle shifts and propose a simple method to reduce its magnitude. Additionally, since matching sections of the speckle patterns in reference and sample image data sets are shifted relative to one another, the assignment of a found signal value to a pixel location is ambiguous. We discuss and compare the two obvious approaches. Finally, we present newly introduced region-of-interest and masking features.

We introduce the UMPA model in section 2. The minimization of the model's cost function (which has changed significantly compared to the previous version) is discussed in section 3. The above-mentioned additions and improvements compared to the previous version are introduced with example measurements in section 4. A brief conclusion and outlook are given in section 5. Further mathematical details and a code example are provided in Supplement 1.

2. Mathematical model of UMPA

UMPA is a template-matching algorithm that compares intensity-modulated patterns through the use of analysis windows. In addition to lateral displacements, local variations in mean intensity and modulation amplitude between the patterns are also detected. The mathematical foundations have previously been described in [33]. We re-introduce the key features of the model here, and also include extensions added by the new implementation. All variables in boldface represent two-dimensional vectors in the image plane.

We assume that a single image $I_0(\mathbf{r})$ [$\mathbf{r} = (r_x, r_y)$] is acquired with a diffuser, but without a sample in the beam (the "reference image"). After introducing the sample, another image $I(\mathbf{r})$ (the "sample image") is acquired. UMPA performs a comparison between I_0 and I for every pixel \mathbf{r} , by centering an analysis window of $(2N + 1) \times (2N + 1)$ pixels ($N = 0, 1, 2, \dots$) on $I(\mathbf{r})$, and seeking the most similar match of this window in the reference image $I_0(\mathbf{r})$. In the absence of attenuation and dark-field, as well as negligible change of wavefront curvature due to the sample, the effect of the sample is purely described by a local transverse translation $\mathbf{u} = (u_x, u_y)$ of the image content. The effect of attenuation is described by a uniform decrease of intensity, whereas dark-field is modeled as a decrease in the modulation amplitude, which we define as the difference of the speckle pattern to its local mean. In short, the effect of the sample is modeled as:

$$I^{(\text{model})}(\mathbf{r}; \mathbf{u}, T, D) = T \{ D [I_0(\mathbf{r} - \mathbf{u}) - \langle I_0 \rangle(\mathbf{r} - \mathbf{u})] + \langle I_0 \rangle(\mathbf{r} - \mathbf{u}) \},$$

$$\langle I_0 \rangle(\mathbf{r}) = \frac{\sum_{w_x=-N}^N \sum_{w_y=-N}^N \Gamma(\mathbf{w}) I_0(\mathbf{r} + \mathbf{w})}{\sum_{w_x=-N}^N \sum_{w_y=-N}^N \Gamma(\mathbf{w})}. \quad (1)$$

Here, $\mathbf{w} = (w_x, w_y)$ parameterizes the summation over the analysis window. The window function Γ (a 2D Hamming window) is used to gradually decrease the weight of pixels towards the edge of the window. A simpler model without dark-field, which does not require the calculation of $\langle I_0 \rangle(\mathbf{r})$, is easily obtained by setting $D = 1$ in the above equation. UMPA attempts to minimize the sum of squared differences (SSD) between the reference speckle images I_0 , modulated according to Eq. (1), and the speckle images I acquired with the sample. Furthermore, information from multiple images with different lateral displacements between sample and diffuser may be combined. Thus, the estimated images $\hat{\mathbf{u}}, \hat{T}, \hat{D}$ are those for which

$$L(\mathbf{r}; \mathbf{u}, T, D) = \sum_{m=1}^M \sum_{w_x, w_y=-N}^N \Gamma(\mathbf{w}) \left[I_m^{(\text{model})}(\mathbf{r} + \mathbf{w} - \mathbf{S}_m; \mathbf{u}, T, D) - I_m(\mathbf{r} + \mathbf{w} - \mathbf{S}_m) \right]^2 \quad (2)$$

is minimized. The subscript m denotes the images with the m -th diffuser-to-sample position. The quantity $\Gamma(\mathbf{w})$ [same as in the definition of $\langle I_0 \rangle(\mathbf{r})$] here serves to smooth the effect which sharp edges in the input images may have on the output signals.

In the sample-stepping technique (introduced and discussed in section 4.1), the sample is displaced between frames instead of the diffuser. This lateral displacement \mathbf{S}_m of the m -th frame must be reverted to ensure that, for each m , the same sample feature contributes to L in any given position \mathbf{r} . For the conventional diffuser-stepping approach, $\mathbf{S}_m = \mathbf{0}$ for all m . The values found in the minimization (see section 3) are then assigned to the position \mathbf{r} :

$$\hat{\mathbf{u}}(\mathbf{r}), \hat{T}(\mathbf{r}), \hat{D}(\mathbf{r}) = \underset{\mathbf{u}, T, D}{\operatorname{argmin}} L(\mathbf{r}; \mathbf{u}, T, D). \quad (3)$$

3. Minimization of the cost function

The cost function of the model introduced above has four variables (three if D is fixed to 1), which suggests that its minimization produces a trajectory in a four-dimensional (or three-dimensional) parameter space. In the form presented above however, the introduced models are only well-defined for integer values of u_x and u_y , whereas T and D are floating-point numbers. Most common optimization algorithms are not suitable for parameter spaces of such a structure.

To solve this problem, we introduce a partial optimization procedure, namely the minimization of the cost function with respect to T and D (or only T if $D = 1$), for a fixed shift \mathbf{u} . Due to the mathematical structure of the UMPA cost function, this step can be performed analytically (calculation details are provided in section 1 of Supplement 1). By performing this step for each estimate of \mathbf{u} , optimal values for T and D (\hat{T} , \hat{D}), and a cost function value minimized with respect to T and D [$\hat{L}(\mathbf{r}; \mathbf{u}) = L(\mathbf{r}; \mathbf{u}, \hat{T}, \hat{D})$] are calculated. The problem is thus reduced to an optimization in \mathbb{Z}^2 (\mathbf{u} -space) for each pixel.

In order to determine the sample shift $\hat{\mathbf{u}}(\mathbf{r})$, the cost function's global minimum, i.e.

$$\hat{\mathbf{u}}(\mathbf{r}) = \underset{\mathbf{u}}{\operatorname{argmin}} \hat{L}^{(T,D)}(\mathbf{r}; \mathbf{u}) \quad \text{or} \quad \hat{\mathbf{u}}(\mathbf{r}) = \underset{\mathbf{u}}{\operatorname{argmin}} \hat{L}^{(T)}(\mathbf{r}; \mathbf{u}) \quad (4)$$

must be determined ($\hat{L}^{(T,D)}$ represents the version of \hat{L} with dark-field, $\hat{L}^{(T)}$ the one without). In the presented software, this minimization is performed in two steps. Since the cost functions are only defined for integer values of u_x and u_y , the first step minimizes $\hat{L}^{(T,D)}$ or $\hat{L}^{(T)}$ only on this grid of integer shifts, as described in subsection 3.1. However, this step alone yields an insufficient precision for the phase signal (u_x and u_y rarely vary by more than one pixel for typical imaging setups), necessitating a sub-pixel interpolation step of the cost function landscape, described in subsection 3.2.

3.1. Discrete minimization step

Starting at $\mathbf{u} = \mathbf{0}$, the optimization procedure in \mathbf{u} -space is done by successive one-dimensional downhill optimization subroutines, performed alternately in the u_x and u_y directions. The one-dimensional optimization compares cost function values of the current estimate with its neighbors and varies the estimate by one pixel per step. Previously calculated cost function values are cached in order to maximize computation speed. An estimate is considered to be a minimum in one dimension if it is lower than both neighbor values. When this happens, an identical minimization procedure in the orthogonal direction begins, starting from the current estimate. This procedure repeats until an estimate is simultaneously a minimum in both directions. We call this value \mathbf{u}_d the discrete minimum.

Since the refined, sub-pixel-precision minimum must lie within ± 1 pixel of the discrete minimum, it is sufficient to interpolate a 1×1 pixel neighborhood for minimization. It is computationally efficient to interpolate in a 1×1 square *between* the points of the integer grid (see highlighted area in Fig. 1(b)), since the interpolated surface in this area is fully defined by the smallest possible number of neighboring cost function values, namely $4 \times 4 = 16$.

As four such areas ("quadrants") are adjacent to the found minimum \mathbf{u}_d , it must be determined ahead of time which one contains the sub-pixel minimum. This is done by comparing the cost

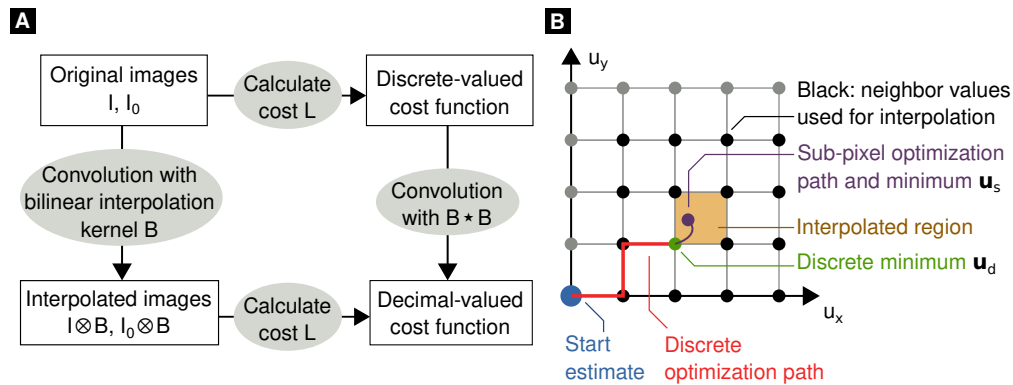


Fig. 1. (a) Approach for cost function interpolation. Performing bilinear interpolation (convolution with the kernel B) on the input images, and then calculating the cost function of the interpolated values is (approximately) equivalent to directly calculating cost function values and then interpolating this with the kernel $B \star B$. (b) Procedure for identifying sub-pixel shift values. From the starting estimate $\mathbf{u} = (0, 0)$, optimization is first done on the grid of integer-valued shifts. Once the discrete minimum \mathbf{u}_d is found, the convolution of the cost function with $B \star B$ is calculated in a 1×1 pixel area adjacent to the minimum. The choice of “quadrant” relative to \mathbf{u}_d is determined by the location of its neighbors with the lowest cost function values. Newton-Raphson optimization on the cost function in this area determines the sub-pixel optimum \mathbf{u}_s . Cost function values from a 4×4 square of pixels surrounding the interpolated area (black dots) are required for the convolution due to the 4×4 pixel footprint of the interpolation kernel. Note that the estimates for T and D are not modified after sub-pixel optimization.

function values of the vertical and horizontal neighbors of \mathbf{u}_d , and selecting the quadrant adjacent to the neighbors with the lowest values.

3.2. Sub-pixel minimization step

A simple method to identify the cost function minimum with sub-pixel precision is to interpolate the recorded images $I_m, I_{0,m}$, and thus evaluate the resulting interpolation functions on a finer grid than the original images. Calculating the cost function from these images then produces a more finely sampled version of the original cost function, thus yielding phase-shift values with greater accuracy. We found that this process can be simplified by directly convolving the original cost function landscape with a specific kernel. This idea is schematically shown in Fig. 1(a).

Bilinear interpolation can be expressed as a convolution of the input function (defined on a grid) with an interpolation kernel B . We show in section 3 of Supplement 1 that the cost function landscape calculated from bilinearly interpolated intensities is given by the convolution of the non-interpolated cost function landscape $\hat{L}(\mathbf{r}; \mathbf{u})$ with the kernel $B \star B$, i.e., the autocorrelation function of B .

Since this kernel has an analytical expression, the interpolated cost function surface and its local first and second positional derivatives do as well. This motivates the use of the Newton-Raphson method to identify the location of the minimum: beginning at the integer minimum \mathbf{u}_d , each update of the current estimate requires calculation of the first and second partial derivatives of the interpolation of \hat{L} . Optimization is terminated when the magnitude of the position update falls below a threshold (10^{-4} pixels by default).

As the used interpolation kernel has a footprint of 4×4 pixels, correct calculation of interpolated costs in a 1×1 pixel square requires evaluation of the cost function in a 4×4 pixel neighborhood around \mathbf{u}_d [see black dots and highlighted area in Fig. 1(b)]. Reusing cached cost function values

from the discrete minimization reduces the number of function evaluations necessary for this. Note also that the estimates for T and D are determined during the discrete minimization step, and are not updated after sub-pixel minimization, as we consider this change negligible.

As illustrated in Fig. 1(b), identifying the sub-pixel minimum of the cost function \hat{L} (i.e., $\hat{L}^{(T)}$ or $\hat{L}^{(T,D)}$) for one image pixel \mathbf{r} is achieved by:

1. Identifying the integer shift \mathbf{u}_d for which \hat{L} is minimized and finding the 1×1 pixel “quadrant” with the lowest neighbors.
2. Evaluating \hat{L} for a 4×4 neighborhood of points (integer shifts) centered on this quadrant
3. Convolution of this neighborhood with the interpolation kernel $B \star B$
4. Newton-Raphson minimization of the interpolated cost function landscape, yielding a sub-pixel minimum \mathbf{u}_s .

4. Improvements by the new UMPA implementation

4.1. Addition of the “sample-stepping” acquisition technique

The UMPA model can combine image data with different relative positions of diffuser and imaged sample (indexed by the subscript m in the preceding equations). There are important benefits for doing so:

- *Improved convergence*: speckle patterns often look similar to many shifted versions of themselves. Thus, if only one diffuser position is used ($M = 1$), the cost function may exhibit several local minima in \mathbf{u} -space, which increases the probability of minimization converging to an incorrect solution. If multiple diffuser positions are combined, however, the full cost function landscape is given by the sum of landscapes for each position. Since the locations of incorrect local minima are mostly uncorrelated for sufficiently large diffuser displacements, they increasingly cancel out, while the correct minimum is reinforced.
- *Spatial resolution*: while UMPA can be used with a single diffuser position ($M = 1$), this is not compatible with the window size parameter $N = 0$ (i.e., an analysis window of 1×1 pixel), since 3 or 4 unknowns would have to be retrieved from only 2 data points. The use of analysis windows with $N > 0$ is however always associated with a decrease in spatial resolution. Thus, pixel-size resolution is only achievable with the use of multiple diffuser steps, an approach similar to the original XSVT model [30].
- *Noise performance*: naturally, combining multiple diffuser positions also increases total detector dose and thus the contrast-to-noise ratio in all retrieved image modalities.

The previously established method to achieve these diffuser-to-sample displacements is to translate the diffuser while the sample is kept stationary. Here we call this approach *diffuser stepping*. However, since only the relative shift between sample and diffuser is crucial, it is also possible to laterally displace the sample, and then modify the mathematical model to take this displacement into account. We call this *sample stepping*. The difference between the two methods is illustrated in Fig. 2, and example datasets acquired with sample stepping are shown in Fig. 3 and Fig. 4. The cost functions are modified to account for sample translations by subtracting the sample displacement \mathbf{S}_m in each frame, as shown in section 2.

Here, \mathbf{S}_m is a 2D vector describing the relative transverse position of the sample in the m -th frame, in multiples of the effective pixel size, and relative to the starting point (e.g., the first frame). For diffuser stepping, \mathbf{S}_m is simply $\mathbf{0}$ for all m . Conversely, the location of the diffuser is not explicitly required for signal retrieval (neither for sample stepping nor for diffuser stepping), and thus not expressed in the mathematical model presented here.




Fig. 2. Illustration of the two available stepping methods. (a) *Diffuser stepping*: the diffuser is laterally translated, while the sample remains stationary. (b) *Sample stepping*: the diffuser remains stationary and the sample is translated. (c) “Coverage maps” of the sample for the examples in (a) and (b), showing how many times each part of the sample is imaged. Sample stepping increases the size of the covered region, at the expense of a locally decreased coverage.




Fig. 3. Flower (Hare’s foot clover) imaged at the P05 beamline, PETRA III, DESY, Hamburg with the sample-stepping technique. (a): Overview of the whole sample (\hat{u}_x). The dashed rectangle represents the size of the detector field of view (3.48 mm \times 7.25 mm). Three regions of interest (solid rectangles) are shown in the insets on the right, for all four modalities (columns from left to right: horizontal shift \hat{u}_x , vertical shift \hat{u}_y , dark-field \hat{D} , transmittance \hat{T}). The images are processed with the window size parameter $N = 2$. The sample was translated on a rectangular grid of 25×10 positions, yielding a measurement time of 27 min. $E = 35$ keV, pixel size: 0.916 μm . Sample-detector distance: 20 cm, diffuser (9 \times 400-grit sandpaper) 11.5 cm upstream of sample. Scale bar (white): 1 mm. The original size of (a) is 12.4 mm \times 23.5 mm, i.e., 13500 \times 25700 pixels, but the shown data is binned [10 \times 10 in (a), 3 \times 3 in (b)–(n)]. UMPA processing time was 6.5 min, using 120 threads.




Fig. 4. Centipede imaged at the SYRMEP beamline (Elettra Sincrotrone, Trieste, Italy) with the sample-stepping technique. (a): transmittance \hat{T} , (b, c): horizontal and vertical refraction \hat{u}_x, \hat{u}_y (in pixels). The dark-field modality is omitted here due to a lack of visible features. The dashed rectangle represents the detector field of view (5.46 mm \times 17.46 mm). The sample was translated on a rectangular grid of $M = 2 \times 26$ positions, for a total measurement time of 12 min. UMPA processing time (with $N = 2$) was 14 s, using 50 threads. $E = 20$ keV, sample-detector distance 32 cm, diffuser (5 \times 320-grit sandpaper) 111.2 cm upstream of sample. Eff. pixel size: 8.8 μ m. Scale bar is 2 mm.

Figure 2(c) schematically shows “coverage maps” achieved with the different stepping approaches, i.e. the map of the number of times each part of the sample was part of an acquisition. Diffuser stepping always yields constant coverage across the field of view (FOV), while the coverage of sample-stepping measurements encompasses a larger area, but gradually decreases towards the edge of this area. This leads to a decrease of image quality (poorer convergence and higher noise levels, especially in the differential-phase and dark-field modalities) in these regions. Besides this basic trade-off, these are some of the main benefits of using the sample stepping approach:

- *Faster acquisition of very large samples:* acquiring an object far larger than the field of view with diffuser stepping generally requires a larger number of motor movements (doing a full diffuser stepping, moving on to the next sample position, etc.), and thus requires more overhead time.
- *High-throughput imaging:* for low- or medium-resolution applications, sample stepping could be performed with continuous sample displacement, e.g., via a conveyor belt. This would eliminate motor displacement overhead and could lead to drastic speed-ups in acquisition time, similar to “fringe-scanning” used in grating-based X-ray [35].
- *Simpler instrumentation:* only one motorized stage is required (the one for the sample). Furthermore, no high repeatability is required for this stage, since each position is

approached only once. An accurate readout of this stage's position is still useful, but deviations can be corrected through cross-correlation analysis.

However, the approach may also introduce some (situational) drawbacks:

- *Drifts*: for longer acquisitions with many frames (e.g., a large two-dimensional grid of sample positions), drifts of sample and/or diffuser positions may occur, leading to artifacts. However, this can be ameliorated by frequently acquiring reference images, and deviations can often be corrected with cross-correlation analyses.
- *Identifying sample limits*: determining the extent of a large object in terms of sample motor positions can be difficult, especially when the reduction of coverage towards the edge of the reconstructed area is taken into account (cf. Fig. 2(c)). The sample may then hardly be visible for the "edge" sample positions. For samples smaller than the detector FOV, care should be taken that the object does not leave the FOV for any of the motor positions.
- *Effective pixel size must be known / determined*: this is necessary since the displacement vectors \mathbf{S}_m must be expressed as multiples of the pixel size. However, performing UMPA reconstruction with different pixel size estimates and comparing the quality of output images yields quick and precise results.
- *Variable image quality / noise levels due to the non-constant coverage*. For achieving roughly constant coverage in the center of the reconstructed regions, we recommend the use of a one- or two-dimensional grid of motor displacements, equal to an integer fraction of the field of view, i.e.: $\Delta x = W_x/N_x$, $\Delta y = W_y/N_y$, where W_x and W_y is the extent of the detector FOV in x and y , and $N_x, N_y \in \mathbb{N}$. This yields a maximum coverage of N_x or N_y for 1D stepping, and $N_x N_y$ for a 2D stepping grid. Furthermore, a helper function is available in the software package to plot coverage maps for a given sample trajectory.
- *Limited compatibility with cone-beam geometry*: in setups with a high ratio of detector FOV and source-to-detector distance, sample stepping may produce artifacts due to projection inconsistencies between different sample displacements, especially for thick samples. (This effect is exploited by digital tomosynthesis to extract depth information.) However, it should be noted that similar effects occur when tiling diffuser-stepping reconstructions, and that the sample-stepping technique is probably less crucial for cone-beam setups, since these can achieve larger fields of view.

4.2. Speedup and software architecture

C++ is a very widely used programming language that combines high computational speed with the ability to use object-oriented programming concepts, thus making it suitable for complex applications where high efficiency is required. Due to its extensive use, C++ compilers exist that can generate highly optimized code for nearly all computing architectures and operating systems. However, these advantages come at a cost: C++ code is often more complex, less readable, and thus more difficult to modify than an equivalent implementation in an interpreted language such as Python. Due to its high flexibility and ease of use, Python has been widely adopted by the scientific computing community, and a large ecosystem of Python modules for numerical and scientific applications has since emerged.

In order to bridge the gap between the worlds of Python and C/C++, the Cython language has been developed [34]. It features a syntax similar to Python, but is compiled into C or C++ code (and then into machine code by a C/C++ compiler). This enables the two main use cases of Cython: writing high-performance code with a readability and flexibility similar to Python, and integrating existing C/C++ modules in a Python interface. Both aspects were used for the work presented here in order to combine the high efficiency of C++ with the accessibility of Python.

Due to the language's implementation, Python programs are usually limited to a single process. This is of crucial importance for UMPA, because its central problem, the minimization of the cost function, is "embarrassingly parallel", i.e., it is easy to parallelize with minimal overhead, since the problem is independent for each image pixel. Parallelization in Python can still be achieved by spawning multiple processes, but this is slower than multithreading, has greater overhead, and is more prone to complications. With Cython however, multiprocessing can be achieved easily and with minimal overhead via OpenMP. The speedup achieved by the use of C++ and multithreaded processing are illustrated in Table 1.

Table 1. Runtime measurements of the Python and C++ implementations of UMPA.^a

Image size [px]	M	N	UMPA version	Threads	Runtime	Rel. speed-up
1000 × 1000	25	2	Python	1	3864 s	1
			C++/Cython	1	(30.90 ± 0.03) s	125
			C++/Cython	16	(2.10 ± 0.02) s	125 × 14.7
			C++/Cython	64	(0.89 ± 0.05) s	125 × 34.7

^aBenchmarks were run on a compute server with 512 GByte of memory and two AMD EPYC 7551 CPUs with 32 cores each. Runtime values for the C++/Cython benchmarks are given as mean and standard deviation of five runs.

The core element of the presented software package is the implementation of the cost function calculation, and its minimization, in C++. A Cython module provides the interface between these C++ libraries and the Python front-end. A class-oriented program structure is used both in the C++ and Cython sections of the module, which avoids repetition of code between the two available UMPA models (with and without dark-field). A code example with an explanation of the individual steps is provided in section 2 of Supplement 1.

As illustrated in Table 1, the new C++ implementation of UMPA running on a single thread is about 125 times faster than the Python implementation, while multithreading adds an additional speed-up factor closely related to the number of used threads.

4.3. Correction of differential-phase estimation bias

The presented sub-pixel interpolation approach introduces a bias to the determined shifts u_x, u_y . This is most obvious in sample-free areas of the FOV, or when comparing a set of images with itself, i.e., running UMPA with $I_{0,m} = I_m$. Although the resulting cost function values for integer shifts have a minimum at $\mathbf{u} = \mathbf{0}$, the determined sub-pixel minimum often has an offset.

On the other hand, explicitly oversampling a set of images ($I_{0,m}, m = 1, 2, \dots$) and calculating the cost function of the image set with itself must still yield a minimum precisely at $\mathbf{u} = \mathbf{0}$: a sufficiently "random" pattern should be most similar to an unshifted copy of itself.

This apparent contradiction is probably resolved by the fact that the interpolation kernel is not quite equivalent to oversampling of the raw data: The interpolation step is exact for the cost function terms I_1, \dots, I_6 , and thus also for any linear combination of these terms, such as $L^{(T)}$ and $L^{(T,D)}$ [see Eqs. (S1), (S2), and (S5) in Supplement 1]. However, the minimization of these cost functions with respect to T , or T and D respectively, introduces non-linear operations (division and multiplication of the I_j terms, see Eqs. (S4), (S7), and (S8) in Supplement 1. These operations do not commute with oversampling, which probably causes the observed bias.

Although the exact origin of the effect has not been fully ascertained, we found that the lateral shifts of the reference data set I_0 with itself is a good estimate for this bias. We thus use the following procedure to reduce the estimation bias of $\hat{\mathbf{u}}$:

1. Calculation of $\hat{\mathbf{u}}(\mathbf{r})$ (including sub-pixel interpolation) with I_m as sample data and $I_{0,m}$ as reference data

2. Repeating the same calculation, but with $I_{0,m}$ as *both sample data and reference data*, i.e. substitution of I_m with $I_{0,m}$, yielding the shifts $\hat{\mathbf{u}}^{(\text{ref})}(\mathbf{r})$
3. Calculation of $\hat{\mathbf{u}}^{(\text{corr})}(\mathbf{r}) = \hat{\mathbf{u}}(\mathbf{r}) - \hat{\mathbf{u}}^{(\text{ref})}(\mathbf{r})$

An example for this correction is shown in Fig. 5. This corrects the majority of the bias for weak refraction (i.e., significantly less than 1 px). For larger amounts, the effectiveness of this correction decreases, although image quality is still improved. We have found that the bias correction remains more accurate if the correction is taken at a different pixel, shifted by $\hat{\mathbf{u}}$:

$$\hat{\mathbf{u}}^{(\text{corr})}(\mathbf{r}) = \hat{\mathbf{u}}(\mathbf{r}) - \hat{\mathbf{u}}^{(\text{ref})}(\mathbf{r} + \hat{\mathbf{u}}). \quad (5)$$




Fig. 5. Illustration of bias correction. (a): horizontal shift \hat{u}_x from UMPA, (b): bias estimate $\hat{u}_x^{(\text{ref})}$ obtained by matching reference images with themselves, (c): difference between (a) and (b). The fixed-pattern noise is clearly reduced, yielding an improved contrast for faint structures (here: zoomed region of the flower shown in Fig. 3).

4.4. Comparison of coordinate assignment schemes

The cost function introduced in section 2 depends on an analysis window on the sample image(s), centered on some coordinate $\mathbf{r}' = \mathbf{r} - \mathbf{S}_m$, and a second window on the reference image(s), shifted by an amount \mathbf{u} , i.e. centered on $\mathbf{r}' - \mathbf{u}$. This notation implies that the reference window moves during the minimization procedure (since \mathbf{u} is the optimized variable), and the sample window remains fixed. In particular, the values returned by the minimization procedure (\hat{T} , \hat{D} , and $\hat{\mathbf{u}}$) must then be assigned to \mathbf{r}' , the center of the sample window. This is illustrated in Fig. 6(b). If the values were instead assigned to the center of the *moving* window, there is a possibility that some pixels of the output images are never being assigned a value, and other pixels being assigned a value multiple times, since the UMPA optimization loop is run exactly once for every value of \mathbf{r}' .

However, Fig. 6(a) illustrates that it is more physically accurate to assign the UMPA output to the center of the reference window, since this point coincides more closely with the point of interaction between radiation and object. This also corresponds with a modification suggested in [22] [Eq. (20)]. Therefore, we have introduced an alternative mode for the optimization procedure where the reference window remains fixed and the sample window is moved during the optimization. This is simply achieved by adding \mathbf{u} to all coordinates, i.e. the reference window is now centered on \mathbf{r}' , and the sample window is centered on $\mathbf{r}' + \mathbf{u}$ (see Fig. 6(c)). In the code, either minimization procedure can be selected with the parameter `assign_coordinates` (see line 10 in the code example, section 2 of Supplement 1). In this case, the UMPA model

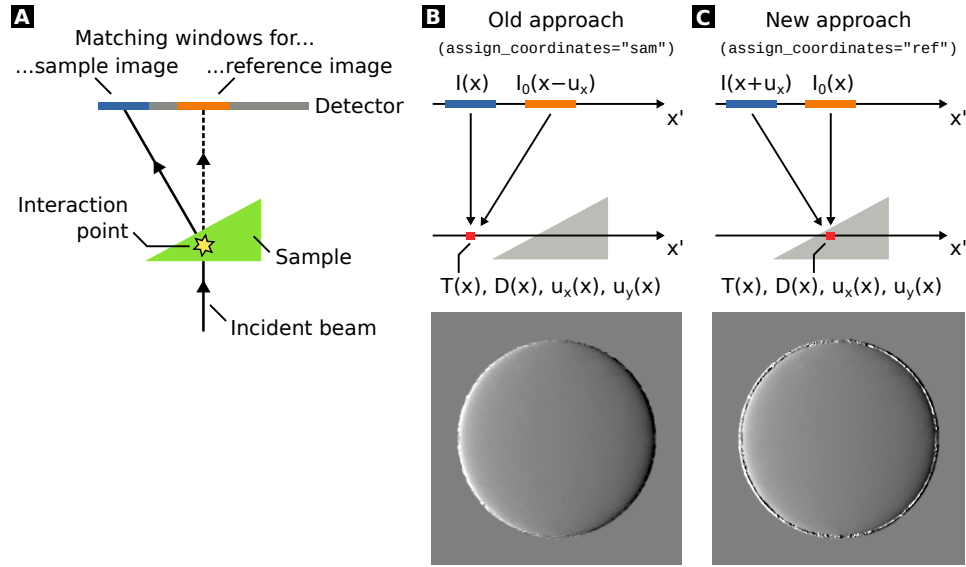


Fig. 6. Illustration of the two coordinate assignment schemes. (a) Refraction by the sample induces a shift between matching analysis window pairs in the sample and reference images, prompting the question of which position the algorithm's output should be assigned to. (b) For the old approach (`assign_coordinates="sam"`), the calculated image data (i.e., transmittance, differential phase and dark-field) is assigned to the sample window center. (c) For the new approach (`assign_coordinates="ref"`), the values are instead assigned to the reference window center. Intuitively, this is more reasonable since the assignment point coincides with the interaction point of sample and incident light, but the shown example image shows that it introduces artifacts near sample edges. We found that this is due to UMPA avoiding sample regions with strong distortions of the speckle pattern.

[Eq. (1) and (2)] changes to:

$$I^{(\text{model})}(\mathbf{r}; T, D) = T \{D [I_0(\mathbf{r}) - \langle I_0 \rangle(\mathbf{r})] + \langle I_0 \rangle(\mathbf{r})\},$$

$$L(\mathbf{r}; \mathbf{u}, T, D) = \sum_{m=1}^M \sum_{w_x, w_y=-N}^N \Gamma(\mathbf{w}) \left[I_m^{(\text{model})}(\mathbf{r} + \mathbf{w} - \mathbf{S}_m, T, D) - I_m(\mathbf{r} + \mathbf{u} + \mathbf{w} - \mathbf{S}_m) \right]^2. \quad (6)$$

However, the example images in Fig. 6 (generated from simulated speckle data of a sphere) show that this new approach performs poorly near sample edges. This is due to a distortion of speckle patterns in the presence of sample edges (e.g., due to strong wavefront curvature or propagation fringes), thus exhibiting low similarity to all regions in the reference speckle pattern.

For the new approach (`assign_coordinates="ref"`, see lines 10 and 14 in the code example), this leads to a cost function increase as the sample window approaches an edge, and the algorithm thus avoids selecting such sample window positions, leading to the black-and-white edge artifact in Fig. 6(c). For the old approach (`assign_coordinates="sam"`) however, the optimization is performed with a stationary sample window. In the case of a sample window centered on an edge, the cost function increase due to the distortion of the pattern is present for all shifts \mathbf{u} (i.e., all possible reference window matches), and thus essentially behaves as an offset to the cost function landscape, having a much weaker impact on the minimization procedure.

The old approach is therefore maintained as the default method in the present implementation. Since the two methods essentially minimize different subsets of the same cost function space, we believe it is possible to synthesize both approaches into a method which is both physically

accurate, and numerically robust. One option might be to process with the old approach, and then correct the misattribution of positions by applying an image warping transform to the images, using a warp map generated from the found shift values $\hat{\mathbf{u}}(\mathbf{r})$ [or the bias-corrected $\hat{\mathbf{u}}^{(\text{corr})}(\mathbf{r})$]. This being said, we believe that the magnitude of this misattribution is minimal for most imaging tasks, since the vast majority of retrieved differential phase shifts are in the vicinity of the setup's resolution limit.

4.5. Addition of options for weighting and region-of-interest processing

The subset of data to include in the UMPA reconstruction can be reduced arbitrarily through the use of a pixel mask (see parameter `mask_list` in the code example, section 2 of Supplement 1). This mask has the same dimensions as the input data, so that the contribution of each raw data pixel can be individually controlled. This is especially useful if only some of the raw image frames in a dataset contain artifacts. A typical use case for masks is with an iterative workflow, i.e.: processing the full data set, identifying image regions containing artifacts, excluding the matching data subsets in the mask, and re-processing. Alternatively, mask values can be set to non-boolean values and will then be interpreted as a weighting for the UMPA cost function.

If sample stepping is used and a region to be excluded is stationary relative to the sample, formulating a suitable mask to exclude it in all frames may be difficult, since this region must be tracked with the sample's movement [cf. Fig. 2(b)]. In this case, the ROI feature can be useful: it functions similarly to the "mask" feature, but coordinates are given in the reference frame of the reconstructed image, not the raw data. The ROI parameter is given as a pair of NumPy slice objects or `(start, stop, step)` pixel index ranges. It can thus be used to reduce the reconstructed field of view, or to reduce the resolution of the reconstruction by calculating only every `step`-th pixel value. The latter is especially useful if very large analysis windows are used: since the resolution of output data is restricted by the window size, the `step` parameter can be safely increased without loss of information.

5. Conclusion and outlook

We presented an improved implementation of the "Unified Modulated Pattern Analysis" algorithm. Crucial implementation details, such as the definition of the cost function and its minimization, were discussed in-depth. The software is available in a public GitHub repository (see the "Data availability" statement below). It includes support for the "sample stepping" technique, which obviates the need for motorized displacement of the structure generating the modulated intensity pattern, and allows imaging samples far larger than the detector field-of-view. We demonstrated the feasibility of this technique with several examples.

As the software is far more computationally efficient than the preceding Python implementation, and is also capable of multithreading, it is ideally suited for processing large volumes of wavefront-marking X-ray imaging data, such as tomography datasets or large-field-of-view sample-stepping projection images. Even the combination of both techniques, "sample-stepping tomography", where the tomography axis is laterally displaced between tomographic scans, is feasible. The presented software has already been successfully used in published work with two-dimensional gratings [15] and sandpaper [36] as wavefront-marking devices.

We discussed the presence of an estimation bias in the produced differential-phase images and showed a simple approach to reduce this effect. Finally, we discussed the inherent ambiguity in the assignment of fit results to pixel coordinates. We compared a newly introduced method to resolve this problem with the older, previously used approach and found that, although the new approach is in principle more physically accurate, it produces image artifacts of greater magnitude near sample edges.

Besides this, the physical model underlying UMPA can be further refined: in particular, the mathematical model used for the dark-field signal is an imperfect description of the underlying

physics, since all spatial frequencies of the speckle pattern (except zero) are scaled by the same factor. Instead, it is more realistic to characterize small-angle scatter by convolution with a spatially variable Gaussian blur kernel. The adaptation of UMPA to such a model, and its successful application to measurements of highly directional carbon fiber samples has recently been demonstrated [36].

If sufficient data read speeds are achieved, additional speedups for UMPA could be achieved by compiling it for GPU execution. As the core components of the presented implementation are written in C++, they could be modified for e.g. CUDA or OpenCL with limited effort. Another recent publication has also explored GPU acceleration for speckle-based X-ray processing [37].

For the use of UMPA in tomographic imaging, as with many X-ray phase-contrast imaging methods, pre-processing of input data can have a strong impact on image quality, especially if the stability of the setup over time is poor (i.e., mechanical instabilities of setup or optics components). Besides common tomographic pre-processing routines (e.g., dark-current, flat-field, and ring-artifact corrections), the identification and correction of drifts of the diffuser (or other wavefront-marking element) is an important step specific to wavefront-marking techniques. Since the differential-phase signal is on the order of one pixel, it is affected even by minute drifts. Cross-correlation techniques can usually determine and correct these with limited computational effort [38]. On the other hand, a correction of sample position drifts, independent from the diffuser drift correction, may also be required (especially for the sample-stepping technique), which can be more challenging. Even when such corrections are applied however, low-frequency artifacts in reconstructed phase tomograms (probably a side-effect of the integration of differential-phase data) often pose a problem for discriminating low-contrast features.

In summary, we believe that the readiness of UMPA for routine phase-contrast imaging is primarily constrained by the availability of automatic data pre-processing routines dedicated to speckle-based imaging, and the knowledge of which effect (and thus, which correction) is responsible for an observed artifact in the tomographic reconstruction. Recent additions to the “Algotom” package [39] for speckle-based X-ray imaging have also addressed this issue.

Another aspect to consider for the routine application of UMPA for tomography is the use of efficient acquisition protocols, and the order in which motors are moved: performing tomography scans with continuous rotation, and changing diffuser positions (or sample positions) between scans, is much faster than performing a full stepping procedure for each angle. Shorter scan times also reduce demands to long-term setup stability.

For lab-based tomography applications, the reduced spatial coherence additionally places a premium on the choice of a good diffuser, and a careful shaping of the source’s photon energy spectrum for maximizing visibility, or similar, intensity-gradient-based quantities [40].

Funding. H2020 European Research Council (866026).

Acknowledgments. We thank Dr. Irene Zanette for helpful discussions. We would like to thank Professor Julia Herzen and Mirko Riedel for the invitation to their beamtime at P05/PETRA III in October 2021, which allowed us to perform the measurement shown in Fig. 3. We acknowledge DESY (Hamburg, Germany), a member of the Helmholtz Association HGF, for the provision of experimental facilities. Parts of this research were carried out at PETRA III and we would like to thank Mirko Riedel and Dr. Felix Beckmann for assistance in using P05. Beamtime was allocated for proposal II-20190765. We acknowledge Elettra Sincrotrone Trieste for providing access to its synchrotron radiation facilities and we thank Dr. Giuliana Tromba and Dr. Adriano Contillo for assistance in using the SYRMEP beamline. Beamtime was allocated for proposal 20210351. This publication is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 866026).

Disclosures. The authors declare no conflicts of interest.

Data availability. The code for the presented software package is available at [41,42]. Data underlying the experimental results presented in this paper are not publicly available at this time, but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

References

1. A. Momose, "Recent Advances in X-ray Phase Imaging," *Jpn. J. Appl. Phys.* **44**(9A), 6355–6367 (2005).
2. A. Snigirev, I. Snigireva, V. Kohn, S. Kuznetsov, and I. Schelokov, "On the possibilities of x-ray phase contrast microimaging by coherent high-energy synchrotron radiation," *Rev. Sci. Instrum.* **66**(12), 5486–5492 (1995).
3. P. Cloetens, W. Ludwig, J. Baruchel, D. V. Dyck, J. V. Landuyt, J. P. Guigay, and M. Schlenker, "Holotomography: Quantitative phase tomography with micrometer resolution using hard synchrotron radiation x rays," *Appl. Phys. Lett.* **75**(19), 2912–2914 (1999).
4. D. Paganin, S. C. Mayo, T. E. Gureyev, P. R. Miller, and S. W. Wilkins, "Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object," *J. Microsc.* **206**(1), 33–40 (2002).
5. T. E. Gureyev, D. M. Paganin, B. Arhatari, S. T. Taba, S. Lewis, P. C. Brennan, and H. M. Quiney, "Dark-field signal extraction in propagation-based phase-contrast imaging," *Phys. Med. Biol.* **65**(21), 215029 (2020).
6. D. Chapman, W. Thomlinson, R. E. Johnston, D. Washburn, E. Pisano, N. Gmür, Z. Zhong, R. Menk, F. Arfelli, and D. Sayers, "Diffraction enhanced x-ray imaging," *Phys. Med. Biol.* **42**(11), 2015–2025 (1997).
7. O. Oltulu, Z. Zhong, M. Hasnah, M. N. Wernick, and D. Chapman, "Extraction of extinction, refraction and absorption properties in diffraction enhanced imaging," *J. Phys. D: Appl. Phys.* **36**(17), 2152–2156 (2003).
8. A. Momose, S. Kawamoto, I. Koyama, Y. Hamaishi, K. Takai, and Y. Suzuki, "Demonstration of X-Ray Talbot Interferometry," *Jpn. J. Appl. Phys.* **42**(Part 2, No. 7B), L866–L868 (2003).
9. F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, "Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources," *Nat. Phys.* **2**(4), 258–261 (2006).
10. F. Pfeiffer, M. Bech, O. Bunk, P. Kraft, E. F. Eikenberry, C. Brönnimann, C. Grünzweig, and C. David, "Hard-X-ray dark-field imaging using a grating interferometer," *Nat. Mater.* **7**(2), 134–137 (2008).
11. A. Olivo, F. Arfelli, G. Cantatore, R. Longo, R. H. Menk, S. Pani, M. Prest, P. Poropat, L. Rigon, G. Tromba, E. Vallazza, and E. Castelli, "An innovative digital imaging set-up allowing a low-dose approach to phase contrast applications in the medical field," *Med. Phys.* **28**(8), 1610–1619 (2001).
12. P. R. T. Munro, K. Ignatyev, R. D. Speller, and A. Olivo, "Phase and absorption retrieval using incoherent X-ray sources," *Proc. Natl. Acad. Sci. U. S. A.* **109**(35), 13922–13927 (2012).
13. M. Endrizzi and A. Olivo, "Absorption, refraction and scattering retrieval with an edge-illumination-based imaging setup," *J. Phys. D: Appl. Phys.* **47**(50), 505102 (2014).
14. K. S. Morgan, D. M. Paganin, and K. K. W. Siu, "Quantitative x-ray phase-contrast imaging using a single grating of comparable pitch to sample feature size," *Opt. Lett.* **36**(1), 55–57 (2011).
15. A. Gustschin, M. Riedel, K. Taphorn, C. Petrich, W. Gottwald, W. Noichl, M. Busse, S. E. Francis, F. Beckmann, J. U. Hammel, J. Moosmann, P. Thibault, and J. Herzen, "High-resolution and sensitivity bi-directional x-ray phase contrast imaging using 2D Talbot array illuminators," *Optica* **8**(12), 1588–1595 (2021).
16. S. Berujon, E. Ziegler, R. Cerbino, and L. Peverini, "Two-Dimensional X-Ray Beam Phase Sensing," *Phys. Rev. Lett.* **108**(15), 158102 (2012).
17. S. Berujon, H. Wang, and K. Sawhney, "X-ray multimodal imaging using a random-phase object," *Phys. Rev. A* **86**(6), 063813 (2012).
18. K. S. Morgan, D. M. Paganin, and K. K. W. Siu, "X-ray phase imaging with a paper analyzer," *Appl. Phys. Lett.* **100**(12), 124102 (2012).
19. R. Cerbino, L. Peverini, M. A. C. Potenza, A. Robert, P. Bösecke, and M. Giglio, "X-ray-scattering information obtained from near-field speckle," *Nat. Phys.* **4**(3), 238–243 (2008).
20. M.-C. Zdora, "State of the Art of X-ray Speckle-Based Phase-Contrast and Dark-Field Imaging," *J. Imaging* **4**(5), 60 (2018).
21. S. Berujon, R. Cojocar, P. Piau, R. Celestre, T. Roth, R. Barrett, and E. Ziegler, "X-ray optics and beam characterization using random modulation: experiments," *J. Synchrotron Radiat.* **27**(2), 293–304 (2020).
22. A. J. Morgan, H. M. Quiney, S. Bajt, and H. N. Chapman, "Ptychographic X-ray speckle tracking," *J. Appl. Crystallogr.* **53**(3), 760–780 (2020).
23. A. J. Morgan, K. T. Murray, M. Prasciolu, H. Fleckenstein, O. Yefanov, P. Villanueva-Perez, V. Mariani, M. Domaracky, M. Kuhn, S. Aplin, I. Mohacsi, M. Messerschmidt, K. Stachnik, Y. Du, A. Burkhart, A. Meents, E. Nazaretski, H. Yan, X. Huang, Y. S. Chu, H. N. Chapman, and S. Bajt, "Ptychographic X-ray speckle tracking with multi-layer Laue lens systems," *J. Appl. Crystallogr.* **53**(4), 927–936 (2020).
24. D. M. Paganin and K. S. Morgan, "X-ray Fokker–Planck equation for paraxial imaging," *Sci. Rep.* **9**(1), 17537 (2019).
25. K. S. Morgan and D. M. Paganin, "Applying the Fokker–Planck equation to grating-based x-ray phase and dark-field imaging," *Sci. Rep.* **9**(1), 17465 (2019).
26. K. M. Pavlov, D. M. Paganin, H. T. Li, S. Berujon, H. Rougé-Labriet, and E. Brun, "X-ray multi-modal intrinsic-speckle-tracking," *J. Opt.* **22**(12), 125604 (2020).
27. K. M. Pavlov, D. M. Paganin, K. S. Morgan, H. T. Li, S. Berujon, L. Quénot, and E. Brun, "Directional dark-field implicit x-ray speckle tracking using an anisotropic-diffusion Fokker-Planck equation," *Phys. Rev. A* **104**(5), 053505 (2021).
28. S. Berujon, R. Cojocar, P. Piau, R. Celestre, T. Roth, R. Barrett, and E. Ziegler, "X-ray optics and beam characterization using random modulation: theory," *J. Synchrotron Radiat.* **27**(2), 284–292 (2020).
29. B. Pan, K. Qian, H. Xie, and A. Asundi, "Two-dimensional digital image correlation for in-plane displacement and strain measurement: a review," *Meas. Sci. Technol.* **20**(6), 062001 (2009).

30. S. Berujon and E. Ziegler, "Near-field speckle-scanning-based x-ray imaging," *Phys. Rev. A* **92**(1), 013837 (2015).
31. S. Berujon and E. Ziegler, "X-ray Multimodal Tomography Using Speckle-Vector Tracking," *Phys. Rev. Appl.* **5**(4), 044014 (2016).
32. I. Zanette, T. Zhou, A. Burvall, U. Lundström, D. Larsson, M. Zdora, P. Thibault, F. Pfeiffer, and H. Hertz, "Speckle-Based X-Ray Phase-Contrast and Dark-Field Imaging with a Laboratory Source," *Phys. Rev. Lett.* **112**(25), 253903 (2014).
33. M.-C. Zdora, P. Thibault, T. Zhou, F. J. Koch, J. Romell, S. Sala, A. Last, C. Rau, and I. Zanette, "X-ray Phase-Contrast Imaging and Metrology through Unified Modulated Pattern Analysis," *Phys. Rev. Lett.* **118**(20), 203903 (2017).
34. S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, "Cython: The Best of Both Worlds," *Comput. Sci. Eng.* **13**(2), 31–39 (2011).
35. C. Kottler, F. Pfeiffer, O. Bunk, C. Grünzweig, and C. David, "Grating interferometer based scanning setup for hard x-ray phase contrast imaging," *Rev. Sci. Instrum.* **78**(4), 043710 (2007).
36. R. Smith, F. De Marco, L. Broche, M.-C. Zdora, N. W. Phillips, R. Boardman, and P. Thibault, "X-ray directional dark-field imaging using Unified Modulated Pattern Analysis," *PLoS One* **17**(8), e0273315 (2022).
37. N. T. Vo, H. Wang, L. Hu, T. Zhou, M.-C. Zdora, H. Deyhle, R. C. Atwood, and M. Drakopoulos, "Practical implementations of speckle-based phase-retrieval methods in Python and GPU for tomography," in *Developments in X-Ray Tomography XIV*, B. Müller and G. Wang, eds. (SPIE, 2022).
38. M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," *Opt. Lett.* **33**(2), 156–158 (2008).
39. N. T. Vo, R. C. Atwood, M. Drakopoulos, and T. Connolley, "Data processing methods and data acquisition for samples larger than the field of view in parallel-beam tomography," *Opt. Express* **29**(12), 17849–17874 (2021).
40. T. Zhou, M.-C. Zdora, I. Zanette, J. Romell, H. M. Hertz, and A. Burvall, "Noise analysis of speckle-based x-ray phase-contrast imaging," *Opt. Lett.* **41**(23), 5490–5493 (2016).
41. F. De Marco, P. Thibault, R. Smith, and S. Savatović, "optimato/UMPA: Initial Release of C++ Version," Zenodo (2022), <https://doi.org/10.5281/zenodo.6984740>.
42. F. De Marco, P. Thibault, R. Smith, and S. Savatović, "Code for the "Unified Modulated Pattern Analysis" (UMPA) method of processing speckle-based X-ray imaging data," Github (2022), <https://github.com/optimato/UMPA>.