

# An automated approach to the alignment of compound refractive lenses

Sean Breckling,<sup>a\*</sup> Bernard Koziowski,<sup>b</sup> Leora Dresselhaus-Marais,<sup>b,c</sup> Arnulfo Gonzalez,<sup>a</sup> Ajanaé Williams,<sup>a</sup> Hugh Simons,<sup>d</sup> Paul Chow<sup>e</sup> and Marylesa Howard<sup>a</sup>

Received 22 June 2021

Accepted 14 April 2022

Edited by A. Bergamaschi, Paul Scherrer Institut, Switzerland

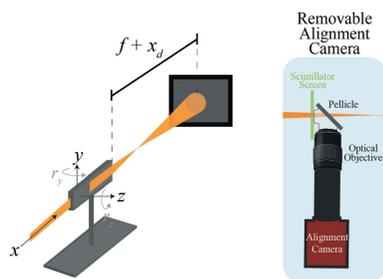
**Keywords:** compound refractive lens; alignment; optimization; beamline optics.

<sup>a</sup>Signal Processing and Applied Mathematics, Nevada National Security Site (NNSS), NV, USA, <sup>b</sup>Lawrence Livermore National Laboratory, Physics Division, Livermore, CA 94550, USA, <sup>c</sup>Department of Materials Science and Engineering, Stanford University, Stanford, CA 94305, USA, <sup>d</sup>Department of Physics, Technical University of Denmark, Fysikvej 311, Kgs Lyngby 2800, Denmark, and <sup>e</sup>HPCAT, X-ray Science Division, Argonne National Laboratory, Argonne, IL 60439, USA. \*Correspondence e-mail: brecklsr@nv.doe.gov

Compound refractive lenses (CRLs) are established X-ray focusing optics, and are used to focus the beam or image the sample in many beamlines at X-ray facilities. While CRLs are quite established, the stack of single lens elements affords a very small numerical aperture because of the thick lens profile, making them far more difficult to align than classical optical lenses that obey the thin-lens approximation. This means that the alignment must be very precise and is highly sensitive to changes to the incident beam, often requiring regular readjustments. Some groups circumvent the full realignment procedure by using engineering controls (*e.g.* mounting optics) that sacrifice some of the beam's focusing precision, *i.e.* spot size, or resolution. While these choices minimize setup time, there are clear disadvantages. This work presents a new automated approach to align CRLs using a simple alignment apparatus that is easy to adapt and install at different types of X-ray experiments or facilities. This approach builds on recent CRL modeling efforts, using an approach based on the Stochastic Nelder–Mead (SNM) simplex method. This method is outlined and its efficacy is demonstrated with numerical simulation that is tested in real experiments conducted at the Advanced Photon Source to confirm its performance with a synchrotron beam. This work provides an opportunity to automate key instrumentation at X-ray facilities.

## 1. Introduction

Lens-based X-ray imaging and diffraction provide opportunities to measure materials with resolutions  $\leq 100$  nm with current technologies. These technologies use either full-field imaging or scanning nano-probes, both of which require X-ray focusing optics. Several X-ray lens options have been developed, including Fresnel zone plates, Kirkpatrick–Baez mirrors and compound refractive lenses (CRLs), which excel for different types of experiments. For imaging applications, CRLs are uniquely versatile, as they have focusing characteristics that can be easily tuned in real time during experiments. CRLs comprise a stack of  $n$  individual lenslets, where the compounded focusing power from each lens element can be changed to tune the effective focal length of the stack. The versatility of CRLs has made them commonplace in synchrotron X-ray experiments, where they are typically used as upstream beam condensers (Vaughan *et al.*, 2011; Schroer *et al.*, 2005) or objective lenses in X-ray microscopes (Lengeler *et al.*, 1999); however, as thick lenses with small numerical apertures, their accurate alignment is



rather challenging. Misalignment of CRLs can increase the width of the focal spot or cause its position to move laterally, resulting in an overall reduction of the quality and intensity of the focal point or image. Corresponding astigmatic aberrations are also generated by the X-ray beam, causing image features to blur directionally as the corresponding focal spot stretches nonuniformly. Together these issues can make the corresponding experiments difficult to interpret.

To align a CRL stack, it must be carefully positioned along the two translation directions ( $y$  and  $z$ ) and two rotational axes ( $r_y$  and  $r_z$ ) that are normal to the axis of propagation of the X-ray beam ( $x$ ), *i.e.* the principal axis of the lens. Optical models to describe CRL alignment (Song *et al.*, 2011) use wavefront simulations to describe the intensity and beam profile as the beam propagates through a CRL with a pre-defined orientation  $\{y, z, r_y, r_z\}$  (positions are typically set by motorized translation/rotation stages). Optimal alignment of CRLs is quite challenging, as the translational and rotational degrees of freedom are coupled, *e.g.* for each position,  $y$ , the total light transmitted varies as the lens is rotated about the  $z$ -axis,  $r_z$ . This condition is specific to CRLs, and arises because of the thickness of the lens, as described fully by Song *et al.* (2011). Simons *et al.* (2017) derived an idealized analytical form for the optical transmission function through CRLs that uses ray-optics to express this coupling. For a given CRL orientation, which we define by the alignment vector  $\mathbf{p} = (y, z, r_y, r_z)$  in the coordinate system described above, the resulting optical transmission function takes the form of a four-variate Gaussian, which we simplify as

$$T(\mathbf{p}) \propto \exp(-\mathbf{p}^T \cdot \mathcal{A} \cdot \mathbf{p}), \quad (1)$$

where  $\mathcal{A}$  is a symmetric  $4 \times 4$  matrix of terms that mostly depend on the lenslet curvatures and number of lens elements in the CRL. Physically, equation (1) derives from the parabolic shape of each lens element and the Beer–Lambert transmission of X-rays through the lens material.

In practice, CRL alignment is performed in two stages: first, a rough alignment is carried out to ensure that some amount of the X-ray beam transmits through the CRL with coarse motorized scans, or guidance from a reference pinhole on the CRL mount. Second, an X-ray detector (either a camera or intensity monitor) is used to finely align the lens based on the transmitted beam. The fine alignment must be reassessed periodically to account for beam drift, sample quality, or other factors over the course of experiments.

A common strategy for fine alignment is to perform two-dimensional motor scans that measure the transmitted intensity as a function of the CRL's position and tilt angles by finding the approximate center of the distribution. Accurate sampling of this search space typically requires detailed scans along each independent axis, often necessitating thousands of individual measurements (samplings) to align a single CRL. A faster method is to carry out several cycles of alternating scans between coupled dimensions; however, this still involves hundreds of samples per iteration. For experiments involving frequent CRL alignment [*e.g.* dark-field X-ray microscopy (Kutsal *et al.*, 2019; Simons *et al.*, 2015)] or that have many

disparate experimental configurations (*e.g.* transmission X-ray microscopy, phase-contrast imaging, or ptychography) the process of aligning, and re-aligning, the CRLs may take a significant fraction of the experimental time.

While experienced researchers may have experiments with minimal drift in the alignment or source, an automated algorithm presents additional benefits to the alignment speed, repeatability, and scalability to multiple simultaneous imaging directions. As X-ray facilities such as synchrotrons and X-ray free-electron lasers (XFELs) often run in continuous operation for weeks or months, a reliable automated method can simplify experimental modifications between users and allow non-expert users to easily realign the experiment if drift or reconfiguration is necessary. This also presents the opportunity in the long term to develop remotely operated or autonomous experimental setups. Automated alignment can also enable facilities to add complexity to their experiments such as multiple probes, smaller incident beam sizes, larger magnifications, or improved spatial resolution depending on the type of experiment. Finally, the repeatability of such an algorithm allows the opportunity to quantify the misalignment and corresponding astigmatism or uncertainty in a measurement – even when substantial jitter or beam-drift is present.

The task of automating the alignment of optics using detection hardware is generally not uncommon, though less common with X-ray lenses at optical facilities. The approach to automating such a task is largely a function of the control hardware, the required accuracy, and the noise levels of the detectors and hardware. One such example is the fine alignment and focus of a camera lens in a wide-field astronomical surveying apparatus. The Robotilter device aligns by optimizing a ‘combo’ score; a positive figure of merit that incorporates both alignment and focus by sampling the four-dimensional search space then performs regression analyses to identify optimal motor positions (Ratzloff *et al.*, 2020). Fang & Savransky (2016) demonstrate the alignment of two independent lens assemblies along the same laser beam. Their procedure searches an eight-dimensional parameter space, determining alignment through a fast study of the spatial distribution of light on a near-field camera sensor. For a given position, the resulting image is decomposed into principal components, which provides feedback for an unscented Kalman filter to determine the next position within the search space to evaluate.

In this work we present an alignment method utilizing a variation of the stochastic Nelder–Mead (SNM) online optimization procedure. Nelder–Mead methods, often referred to as simplex methods, are a class of direct search algorithms that iteratively seek to identify parameters that minimize a prescribed real-valued, nonlinear cost function. In our settings, however, we discuss optimization in terms of maximizing a figure of merit (FOM). The classical formulation of Nelder–Mead presumes an unconstrained search space, and a cost function that is strictly convex. These conditions guarantee that a single ‘best’ solution exists, and is unique. If this condition is not met, *e.g.* because of noise, classical Nelder–Mead often fails to locate an optimal solution. In the litera-

ture, SNM overcomes these difficulties by mollifying noise through increased sampling, and the occasional use of an adaptive randomized search technique when the path toward an optimal solution is not obvious (Lagarias *et al.*, 1998). Our implementation slightly deviates from the literature, opting instead to use a procedural local search based on the Sobol sampling technique (Fox, 1986).

Our description of our method is as follows. We begin in Section 2 by introducing our approach, describing the necessary alignment hardware, and assigning mathematical notation for our methods. In Section 3 we assemble the FOM, which evaluates images from an alignment camera, and returns a transmission estimate akin to (1). We then study this FOM by scanning a 4D space within the limits of motorized stages that position the CRL. These scans are used to provide context for choices made in our algorithm, which is also contained therein. Section 4 reports testing with a series of executions of the algorithm on our apparatus at Advanced Photon Source (APS) (Argonne National Laboratory, USA) as well as several numerical simulations based on data collected there. From the results of collected data, and the numerical simulations, we present a case for improvements and further study of the method in Section 5, and conclude in Section 6.

## 2. Experiment apparatus and mathematical settings

Our apparatus includes the motorized stages that control the position and orientation of the CRL, as well as a simple, self-contained detector assembly (scintillator crystal, imaged onto a camera). We note that the simplicity of the detector and acquisition modes utilized in this work simplifies its implementation in a wide range of experimental settings, allowing versatile and easily customizable configurations. We provide details of our hardware settings and implementation below.

We formulated and tested this method at the Sector 16-ID-D hutch at the APS. This beamline was configured with a double-crystal Si (111) monochromator, and aperture slits to the size of the beam entering the CRL assembly with no additional upstream focusing optics. As such, the beam was nearly collimated and monochromatic with  $\Delta E/E$  of  $1.0 \times 10^{-4}$ . We operated with the monochromator tuned to 9 keV but with the second crystal in the Si (111) monochromator mis-aligned to preferentially remove the third harmonic.

We used a CRL comprising 30 2D Be lenses, each of which has parabolic radii of curvature of  $R = 50 \mu\text{m}$  and a nominal frame thickness of 2 mm, as purchased from RXOptics GmbH (Lengeler *et al.*, 1999) with an effective focal length of  $f = 205 \text{ mm}$ . The lenses were held in a custom-built V-block and purged with nitrogen to prevent damage via oxidation. The full width at half-maximum (FWHM) of the intensity profile immediately downstream of the CRL stack was 0.30 mm. Assuming ray optics, the 0.30 mm aperture and 60 mm length of the CRL stack dictate that the CRL tilt must be within  $\pm 2.5 \text{ mrad}$  ( $0.14^\circ$ ) from the beam axis for some of the collimated incident beam to transmit unobstructed. The CRL assembly itself is mounted on a stage stack, where we make

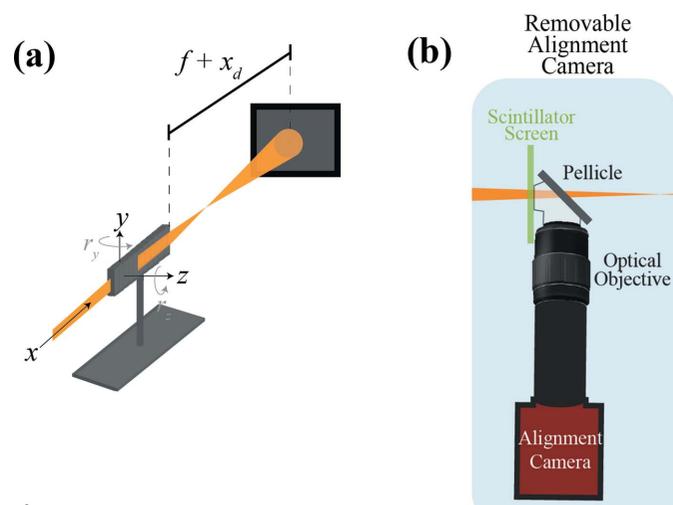
use of four degrees of freedom (two translational and two rotational) in our alignment procedure. We formally denote the translation axes as  $y$  and  $z$ , and the rotations about those axes as  $r_y$ , and  $r_z$ , respectively; thus, the incident beam axis is defined as  $x$ . Given that each motor's position is bounded by travel limits, which are set through software to define the range of the stage's travel, we denote each limiting interval as  $\mathcal{I}_{[i]}$  and define our full search space (the set of all permissible CRL orientations) as  $\Omega$ , the 4D hyper-rectangle of those motor limits is

$$\Omega := \mathcal{I}_y \times \mathcal{I}_z \times \mathcal{I}_{r_y} \times \mathcal{I}_{r_z}.$$

Alignment images were collected in the near field with a YAG scintillator screen of 10 mm diameter and  $50 \mu\text{m}$  thickness, which was placed 200 mm from the CRL, slightly before the X-ray focus. The YAG scintillator absorbs 65% of the incident X-rays, and a pellicle beam-splitter (ThorLabs BP145B1) was used as a dichroic mirror to separate the X-ray and optical beams. A  $7.5\times$  Mitutoyo long-working-distance objective was then used to image the fluorescence from the scintillator crystal onto a Prosilica GC1380 CCD camera. We include a schematic of the apparatus in Fig. 1.

We note that our implementation used the 2D CCD because this allowed the operator the most flexibility in performing the initial alignment and determining best focus manually. It is important to note that spatially resolved intensity is not required for this algorithm and an energy monitor, such as a photo-diode, could be used in place of the imaging detector, as long as it captures the full transmitted beam.

In our formalism, the beam intensity fluctuates in time, and the exposure time of the camera is user-selected, so we define  $t$  as the clock time, and  $\Delta t$  as the camera's exposure time. For a given motor position  $\mathbf{p} \in \Omega$ , and distance from the camera to the CRL focal plane along the beam axis  $x_d$ , we express the raw image collected from the near-field CCD camera as



**Figure 1** Schematic of the basic layout for the alignment detector, detailed in the blue box labeled Removable Alignment Camera, showing how it fits into a synchrotron experiment. We note that most of the detector optics can be customized based on the availability and beam parameters to set the magnification and field of view for the alignment system.

$J_{i,j}(x_d, \Delta t, t, \mathbf{p})$ , where  $i$  and  $j$  represent pixel indices. An upstream ion chamber records a number proportional to the number of photons  $\text{s}^{-1}$  incident on the CRL, which is integrated over  $\Delta t$ . We denote the recorded incident beam count-rate as  $\omega(\Delta t, t)$ , and use it to normalize the fluctuations inherent to the source from our measured transmission, *i.e.*

$$I_{i,j}(x_d, \Delta t, t, \mathbf{p}) := \frac{J_{i,j}(x_d, \Delta t, t, \mathbf{p})}{\omega(\Delta t, t)}. \quad (2)$$

Our intensity measurements do not vary significantly with respect to  $x_d$ . Additionally, normalizing with the beam energy  $\omega(\Delta t, t)$  substantially reduces intensity variation in time  $t$ . Given that we maintain a fixed exposure time  $\Delta t$ , in an effort to aid in readability, we simplify our notation in (2) to  $I_{i,j}(\mathbf{p})$ , or  $I_{i,j}$  when referring to individual pixels on the normalized image, or simply  $I(\mathbf{p})$  or  $I$  when discussing the entire image is prudent and convenient.

The initial rough alignment of the CRL to the beam was performed manually. First, the CRLs were removed completely from the beamline and the position of the X-ray beam on the scintillator was recorded. Next, the CRLs were replaced. The CRL is considered aligned to the X-ray beam when the spot centroid after the CRL falls on the same detector location as when the CRL is not used. Additionally, limits to translational and rotational degrees of freedom were determined by locating the edges of the CRL as shadows in the X-ray beam. The algorithm presented below was tested once the beam and CRL positions were located.

### 3. Building and solving the optimization problem

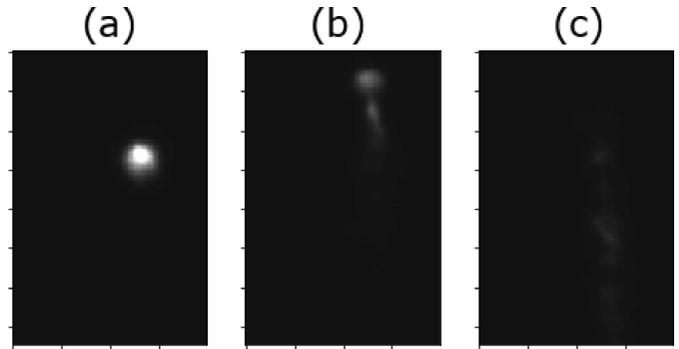
In this section we use the notation from Section 2 to construct an FOM that, when evaluated on our search space  $\Omega$ , behaves sufficiently convex in the presence of standard synchrotron noise. It is the apparent convexity and noise that motivated our choice of the SNM direct-search algorithm, which is presented in a modified form below in Algorithm A.1 in Appendix A.

#### 3.1. Camera-based FOM

The function we develop below is constructed to mathematically discriminate well alignment from misalignment of the CRL using an imaging sensor placed on the beam. When the CRL is well aligned, we expect to see a small, intense spot at the detector plane. When less well aligned, we should see less intense light spread across a larger region of the camera's sensor. Examples are presented in Fig. 2. Here, we develop a robust function that, for a given position  $\mathbf{p}$ , considers the image  $I(\mathbf{p})$  and returns a positive number that quantifies the alignment quality for this system.

For a given motor position  $\mathbf{p}$ , let  $\xi(I)$  and  $\sigma(I)$  denote the median and standard deviation of the resulting detector image  $I(\mathbf{p})$ , respectively. With these terms, we identify a region of interest (ROI) within the image as

$$\hat{I}_M := \{I_{i,j} \in I \mid |I_{i,j} - \xi(I)| > M \times \sigma(I)\},$$



**Figure 2**

Three examples of positions in the search space with the corresponding conditions for alignment and noise issues. Image (a) shows the spot from a well aligned CRL, producing a tight and bright spot, while (b) shows a diffuse tail demonstrating its misalignment, and (c) shows only a diffuse streak from worse alignment still.

where  $M$  is a positive, user-selected threshold parameter. We define our FOM as the median value of the pixels within the ROI, *i.e.*

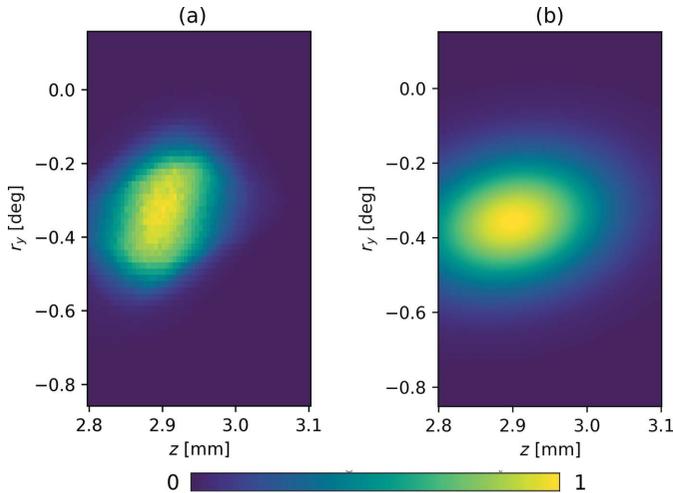
$$F(\mathbf{p}; M) := \xi[\hat{I}_M(\mathbf{p})]. \quad (3)$$

In our apparatus, for any given CRL alignment  $\mathbf{p}$ , the total number of pixels illuminated on the detector plane is substantially lower than half the total number of pixels available. For a well selected threshold  $M$ , the ROI effectively selects, exclusively, for only those illuminated pixels. Therefore, by selecting the median value of the ROI, our FOM (3) returns a larger number for Fig. 2(a) than either Figs. 2(b) or 2(c).

To evaluate the behavior of (3), and to ensure that our model appropriately captures the details of the system, we performed samplings of the full 4D search space  $\Omega$ . While a full, high-resolution 4D scan of  $\Omega$  would have been ideal, given the time-constants inherent to the beamline experiments required we limit our scans to: (1) two independent, high-resolution, uniform scans in 2D, and (2) one lower-resolution quasi-uniform scan in 4D.

To capture the character of  $F(\Omega)$  in 2D raster scans, we first manually aligned the CRL to establish a ground truth well aligned position, which we shall define as  $\mathbf{p}^* = (y^*, z^*, r_y^*, r_z^*)$ . From this position, we performed two raster scans of the  $y \times r_z$  and  $z \times r_y$  axes, fixing the two off-axis parameters to their optimal coordinates, as defined in  $\mathbf{p}^*$ . Using  $\mathbf{p}^*$  as a central position, we selected sampling intervals for each independent axis. Each interval was then discretized as 60 uniformly distributed positions. We present one of the resulting 2D,  $60 \times 60$ -sample raster scan plots in Fig. 3(a).

Collecting a 4D raster scan of  $F(\mathbf{p}; M = 2)$  at the spatial resolution seen in Fig. 3(a) would require an impractical amount of time. In order to collect as much 4D data during the window of time available, we used a space-filling Sobol sequence of positions in  $\Omega$ , as implemented in the `csobol_seq` Python package (Fox, 1986). Given that a Sobol scan does not conform to a grid, we perform a Gaussian


**Figure 3**

Here we directly compare (a) a high-resolution 2D raster scan of the FOM intensity  $F(\mathbf{p}, M = 2)$  and (b) a Gaussian regression of a lower-resolution 4D raster scan. Both scans depict a feature-scaled scan over  $(z, r_y)$ , fixing the remaining variables at  $(y^*, r_z^*)$ .

regression on the results from which a pixel grid can be evaluated. This regression is implemented in the form

$$T_{\text{Gauss}}(\mathbf{p}) = a \exp[-(\mathbf{p} - \hat{\mathbf{p}})^T \cdot \mathcal{A} \cdot (\mathbf{p} - \hat{\mathbf{p}})] + b, \quad (4)$$

where  $a$  and  $b$  are real-valued scalars,  $\hat{\mathbf{p}}$  is the position of the peak, and  $\mathcal{A}$  is a symmetric  $4 \times 4$  matrix – a total of 16 variables. One corresponding 2D slice from this regression is depicted in Fig. 3(b).

### 3.2. Optimizing with SNM

In the absence of noise we would formulate the optimization problem in terms of the idealized model (1) such that the optimal solution  $\mathbf{p}^* \in \Omega$  is given by

$$\mathbf{p}^* := \operatorname{argmax}_{\mathbf{p} \in \Omega} T(\mathbf{p}).$$

However, since synchrotron experiments are always subject to some type of noise, we constructed our FOM  $F[I(\mathbf{p})]$  such that the expectation  $E\{F[I(\mathbf{p})]\}$  is a convex function. Based on this assumption, we formulate the problem in stochastic settings such that

$$\mathbf{p}^* := \operatorname{argmax}_{\mathbf{p} \in \Omega} E\{F[I(\mathbf{p}); M]\}. \quad (5)$$

As long as the character of the noise does not cause the FOM to exhibit errant, problematic local maxima, a direct search method is well suited to solve (5).

Nelder–Mead is a direct-search method that is widely used in nonlinear programming problems where derivatives of the associated cost function are not readily available (Nelder & Mead, 1965; Wright, 1996). The method works by keeping track of a sample set (or simplex) of  $n + 1$  points within the search space, where  $n$  is the spatial dimension of your search space. The simplex is sorted from ‘best’ to ‘worst’, and proceeds by systematically searching for better positions within the search space. New positions are selected as a function of the spatial orientation of the simplex points, and

three real, positive parameters:  $\alpha$ ,  $\gamma$ , and  $\beta$ , known, respectively, as the reflection, expansion, and contraction values. An implementation of the classic method is packaged within a number of popular scientific computing platforms, namely as `fmin` within `scipy.optimize` (Virtanen *et al.*, 2020).

Direct-search methods of this variety are rarely robust to noise. While our FOM (3) is effective at mitigating a substantial degree of the latent measurement noise, it is not entirely eliminated. To account for this, we implemented a modified variety of the SNM method (Chang, 2012; Li & Zhan, 2014).

The algorithm we present below deviates from the literature in four major ways. First, we fix the simplex size to 5, since our search space  $\Omega$  has four dimensions. Second, we formulate our optimization problem to seek a maximal transmission from the detector. The third and fourth modifications both simplify and expedite the FOM sampling criteria seen in the classic versions of SNM, and are detailed below. We enumerate each major step of the SNM procedure with the index  $k$ , beginning with  $k = 0$ . For each  $k$ , when the procedure calls for the FOM to be evaluated, we average  $N_k$  evaluations of (3), where

$$N_k = \max\{\lfloor \sqrt{k} \rfloor, 2\},$$

and  $\lfloor \cdot \rfloor$  returns the nearest lower integer. For a particular position  $\mathbf{p}$ , sample count  $N_k$ , and noise-floor threshold  $M$ , we thus describe the sample-averaged FOM as

$$\mathcal{F}(\mathbf{p}, M, k) := \frac{1}{N_k} \sum_{i=1}^{N_k} F[I^i(\mathbf{p}); M], \quad (6)$$

where  $I^i$  is the  $i$ th image returned from the detector. We determined that the minimum sample of 2 was sufficient, given the noise levels present within the data in our characterization of (3) in Section 3.1. These additional evaluations of (3) come with a minimal time penalty, given that the largest contributor to the wait time per major step  $k$  is the time required to move the motorized stages.

Our main deviation from the SNM formulation described in the literature is our local Sobol search procedure for situations where simplex contraction fails (*e.g.* the final contingency step in Algorithm A.1). In the non-stochastic formulation of Nelder–Mead, this situation is equivalent to deducing that the simplex spans a spatial region too large to effectively continue improving. There, the algorithm performs a ‘shrinkage’ step, which selectively shrinks the region while keeping the ‘best’ positions. In the stochastic formulations seen in the literature, in lieu of shrinking the simplex, a randomized local or global search is performed (Chang, 2012; Li & Zhan, 2014). In these instances, positions within the search space are sampled at random until a new position is found that improves upon the current simplex. In such situations, these SNM formulations utilize an adaptive random search procedure, which potentially performs a global random search of  $\Omega$ . Our proposed alternative to the randomized search procedure is presented in Algorithm A.2.

This deviation was necessary for three reasons. First, while a global random search is a prudent technique to eventually

guarantee mathematical convergence of the SNM procedure on paper, in our use case it comes with a potentially unacceptable time penalty. We first note that the global limits of the motorized stages are substantially larger than the aspect ratio depicted in Fig. 3, thus a full 4D scan is very unlikely to yield successful improvements. Second, note that we have a strong understanding of the character of our FOM  $F(\mathbf{p}, M)$ : a solitary, convex peak that is sufficiently wide in four dimensions. It is unlikely that this FOM would present errant peaks on a similarly configured apparatus. Finally, after a rough alignment of the optics, enough light is present on the detector to initialize the algorithm. The net consequence is that, in these settings, SNM is unlikely to stray far from a very direct path to an acceptable alignment, precluding the necessity of a global search or highly randomized search.

The final step of any successful execution of an optimization algorithm is the stopping condition. We do not formally specify such a condition in our presentation below, but good choices in our use case could be based on the diameter or volume of the current simplex, the total variation of the FOM on the current simplex, by utilizing a measurement from a far-field detector, or by simply enforcing a maximal iteration tally.

### 3.3. Local Sobol search

Sobol sequences are a low-discrepancy, space-filling sampling technique (Sobol, 1967). When each point in a low-discrepancy, space-filling sampling is enumerated, and ordered sufficiently well, we find it to be an efficient method to search a bounded region. In our implementations presented below, we find our Sobol sequences using the `csobol_seq` Python package (Fox, 1986).

To effectively perform a local Sobol sampling during an execution of Algorithm A.1, we need to first identify a suitably local region to explore. Given the current motor position  $\mathbf{p}$  we simply select a rectangular region containing  $\mathbf{p}$  determined by two parameters:  $k$ , the current SNM iteration index, and  $\epsilon$  a positive (small) cooling parameter. Our 4D rectangular region  $\hat{\Omega}$  is then determined to be the region between the two following corner positions,

$$\mathbf{p}_+ = \mathbf{p} + (1 + \epsilon)^{-k} \mathbf{d}, \quad (7)$$

$$\mathbf{p}_- = \mathbf{p} - (1 + \epsilon)^{-k} \mathbf{d}. \quad (8)$$

Some care is required when selecting  $\mathbf{d}$  and  $\epsilon$ . If  $\mathbf{d}$  and  $\epsilon$  are poorly selected, users run the risk of selecting regions too large or too small to search effectively. We found in our implementation that  $\mathbf{d} = d(1.0, 1.0, 1.0, 1.0)$ , where  $d$  is the longest diameter of the initial simplex to be an effective choice, when  $\epsilon$  is selected such that  $(1 + \epsilon)^{-k_{\max}} \simeq 1/10$ , where  $k_{\max}$  is the largest index expected to occur.

When Algorithm A.2 is called from within Algorithm A.1, the local search space  $\hat{\Omega}$  is immediately established by (7) and (8). We then identify an  $N$ -point Sobol sampling of  $\hat{\Omega}$ , which we denote as  $\mathcal{S} = \{\mathbf{p}_i^s\}_{i=1}^N$ . With the motor positions  $\mathbf{p}_i$  established, an expeditious route visiting each position is identified using the nearest-neighbor method, discussed further below.

Selecting  $N$  too large creates a situation where the local Sobol search might take a long time to arrive at a region within  $\hat{\Omega}$  containing an improvement to SNM simplex. Select  $N$  too small, and the Sobol procedure would likely need to be repeated several times, also causing unnecessary delays. For our implementation discussed below, we found choices of  $N$  varying between 10 and 50 to be an effective range.

*Remark 3.1.* A Python implementation and demonstration of Algorithm A.2 can be found online (Breckling, 2022).

Practical implementations of Algorithm A.2 require an efficient way to visit each point in the sample space, given the cumulative time required to move the stepper motors. Let  $\mathcal{R}$  denote a particular course visiting each position within the finite list of coordinates  $\{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\} = \mathcal{S}$ . The time-cost required to move from any one position  $\mathbf{p}_i$  to another  $\mathbf{p}_j$  can be estimated by the Euclidean distance  $|\chi_i - \chi_j|$ . The optimal route  $\hat{\mathcal{R}}$  through  $\mathcal{S}$  minimizes the total Euclidean distance required to travel to each position. While optimal solutions to this variant of the ‘traveling salesman problem’ cannot be practically identified, a nearest-neighbor solution can, and sufficiently approximates  $\hat{\mathcal{R}}$  for our use case.

## 4. Implementation results

In this section we present the results of a study to gauge the efficacy of our automated alignment procedure. In Section 4.1 we highlight the results of a single successful alignment outcome. In Section 4.2 we discuss a collection of results, wherein we attempt to quantify the repeatability of Algorithm A.1.

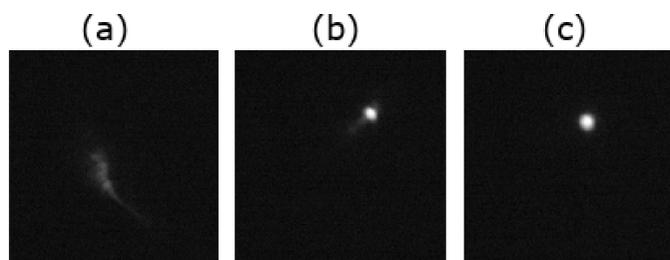
### 4.1. Results of a single alignment procedure

In this section we take a qualitative look at the results of a single execution of Algorithm A.1. Upon manually finding a position  $\mathbf{p}_0$  such that light is visible on the near-field sensor, we randomly generate an initial five-position simplex  $\chi_0$ . This is done by generating random four-vectors,  $\mathbf{r}_k = \{r_i\}_{i=1}^4$ , where each term is sampled from the a uniform distribution on the interval  $[-0.05, 0.05]$  (in mm or  $^\circ$ , where appropriate), then adding such that

$$\mathbf{p}_k = \mathbf{p}_0 + \mathbf{r}_k$$

for  $k = 1, \dots, 4$ . The sampling interval was selected *ad hoc*, but was large enough to see a substantial perturbation of light present on the near-field camera in each of the four dimensions.

The initial simplex  $\chi_0$  is then used to initialize Algorithm A.1. During initialization, the FOM (6) is evaluated for each position within the initial simplex. We select the Sobol sampling count  $N$  to be 10, the cooling parameter  $\epsilon = 2.0 \times 10^{-2}$ , and the local search region to be determined by the vector  $\mathbf{d} = 2.5 \times 10^{-2}(1.0, 1.0, 1.0, 1.0)$  (in mm or  $^\circ$ , where appropriate).



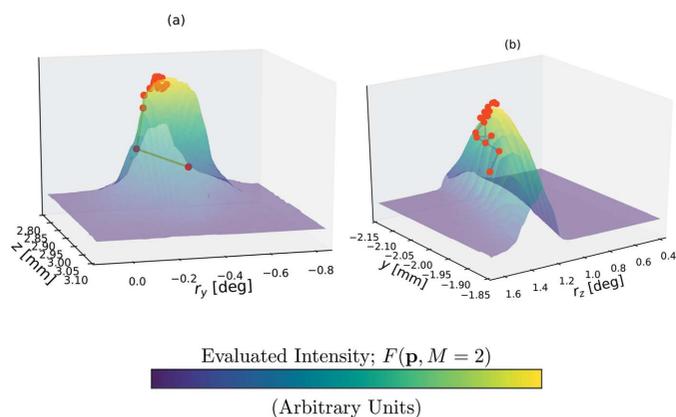
**Figure 4**  
Panels (a), (b), and (c) depict  $I_0$ ,  $I_6$ , and  $I_{62}$ , respectively, which correspond to  $\mathbf{p}_0$ ,  $\mathbf{p}_6$ , and  $\mathbf{p}_{62}$ , the initial, sixth, and final ‘best’ position recorded during one execution of Algorithm A.1. Evaluating each image with the FOM (3), we see that  $\mathcal{F}(\mathbf{p}_0, M)$  and  $\mathcal{F}(\mathbf{p}_6, M)$  correspond, respectively, to 14.03% and 48.29% of  $\mathcal{F}(\mathbf{p}_{62}, M)$  when  $M = 2.0$ .

Upon initializing Algorithm A.1, each position in  $\chi_0$  is evaluated and sorted. The initial ‘best’ position is denoted as  $\mathbf{p}_0$ , with the corresponding near-field camera image  $I_0$ . As the algorithm proceeded to ascend the FOM, we recorded each step, and present the sixth and final ‘best’ positions as  $\mathbf{p}_6$  and  $\mathbf{p}_{62}$  with images  $I_6$  and  $I_{62}$ , respectively, in Fig. 4.

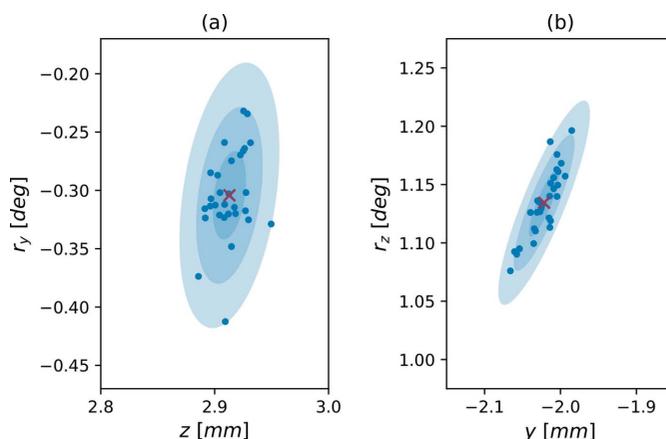
This execution ceased to present any apparent improvements upon arriving at the position  $\mathbf{p}_{62}$ . A signal to end the procedure was then sent. This particular ascent taken by Algorithm A.1 was, with rare exception, a direct one. From initialization to quit, the full procedure required only 62 evaluations of  $\mathcal{F}$ . The incremental improvements seen at most steps appear to be conservatively small. Of those 62 evaluations, there were 16 ‘new bests’, which are presented in Fig. 5.

#### 4.2. A repeatability study

Given that our proposed alignment procedure is inherently stochastic, we performed a Monte Carlo study to evaluate its reproducibility, initializing and executing Algorithm A.1 to a satisfactory alignment 30 times. Our goal was to quantify the spread of the final ‘best’ motor positions. Each execution was initialized in the same way as was done in Section 4.1, *i.e.* given



**Figure 5**  
We present the route taken by Algorithm A.1 ascending the FOM (6). We overlaid the 16 sequential ‘best’ positions upon a surface plot of the raster scans. Panel (a) depicts the  $z$  and  $r_y$  axes, and the  $y$  and  $r_z$  axes can be seen in panel (b). The opacity of the surface was decreased so the full paths remain visible.



**Figure 6**  
Together, these scatter plots depict the best motor position recorded prior to manual termination of 30 executions of Algorithm A.1 (blue dots). The cross (red) represents the average of these recorded positions. We further illustrate the 68, 95, and 99.7% uncertainty regions as concentric ellipses centered at the mean. Panel (a) presents the  $z$  and  $r_y$  axes, while the  $y$  and  $r_z$  axes are given in panel (b). Each execution was initialized from the same region in  $\Omega$ , but the initial simplex was selected randomly (distributed uniformly) about  $\mathbf{p}_0$ .

a fixed motor position resulting from a rough initial alignment  $\mathbf{p}_0$ , complete the initial simplex  $\chi_0$  by sampling randomly within the same fixed rectangular region about  $\mathbf{p}_0$ . In each case, the execution was halted when the resulting spot visible on the alignment camera was sufficiently symmetric and intense.

Over the 30 executions of Algorithm A.1, the average tally of evaluations of the sampling-averaged FOM  $\mathcal{F}$  is 63.67, with a median tally of 62, and standard deviation of 4.84 evaluations. This tally is not to be confused with the index  $k$ , which enumerates the steps within Algorithm A.1. The total number of evaluations of  $\mathcal{F}$  varies from step to step, depending on the location of the simplex. The total-evaluations tally therefore better represents the total time required to fully execute the procedure.

Upon sending the stopping command to end the execution, the final ‘best’ position was recorded. Scatter plots of these ‘best’ results are seen in Fig. 6. We used the spatial statistics of the ‘best’ positions recorded in 30 executions of Algorithm A.1 to generate the two-dimensional projections of the 4D elliptical uncertainty regions presented in Fig. 6. We found that, when the algorithm is manually terminated, 95% of the ‘best’ transmission measurement results were within 90% of the ground truth intensity.

#### 4.3. Simulated repeatability study

The principal goals of our simulated studies are to provide validation for the results presented in the previous section. Second, we intend to demonstrate the efficacy of Algorithm A.1 in terms of noise. We begin by developing a mathematical model of the FOM (3), accounting for estimates of beam jitter and a parameterizable additive noise term. We then test the performance of Algorithm A.1 using the same settings, while prescribing varying amounts of additive noise.

The modeled FOM  $\mathcal{F}_{\text{Gauss}}$  (9) utilizes the 4D Gaussian fit,  $T_{\text{Gauss}}$  (4), and includes a normally distributed additive noise term,  $\epsilon \in \mathcal{N}(0, \zeta)$ , where  $\zeta$  is a user-selected parameter. A baseline estimate of  $\zeta$  is computed as  $5.12 \times 10^{-3}$ . This was accomplished by selecting a region of  $\Omega$  far away from the positions corresponding to well alignment. Additionally, given a beam divergence of  $6.5 \times 10^{-3}$  rad, we consider the possibility of random pointing jitter in the beam itself. Thus, we include normally distributed random perturbations on the  $r_y$  and  $r_z$  axes with a standard deviation of 10% of the beam divergence.

For each Monte Carlo study, we execute Algorithm A.1 a total of 30 times, using the model FOM,

$$\mathcal{F}_{\text{Gauss}}(\mathbf{p}) = T_{\text{Gauss}}^*(\mathbf{p} + \Delta\theta) + \epsilon, \quad (9)$$

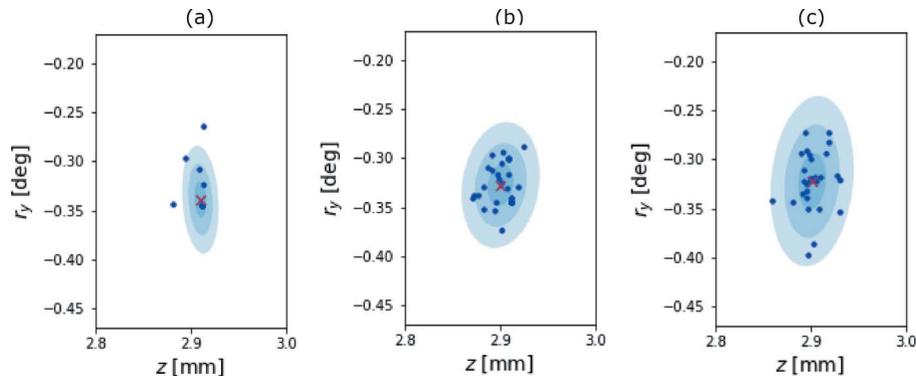
where  $T_{\text{Gauss}}^*$  is unit normalized from  $T_{\text{Gauss}}$ . Each initial simplex is generated around a fixed position  $\mathbf{p}_0$ , then populated by selecting nearby positions at random, as was done in the beamline study in Section 4.2. The number of Sobol sampling positions  $N$  remains set to 10, as with the cooling parameter  $\epsilon = 2.0 \times 10^{-2}$ .

Unlike the synchrotron implementation, which was manually exited by the user, an automatic stopping condition within Algorithm A.1 is required. In the synchrotron study, on average, 63.67 evaluations of (6) were required. In an effort to make a fair comparison, we forced the automated method to exit upon reaching 64 evaluations.

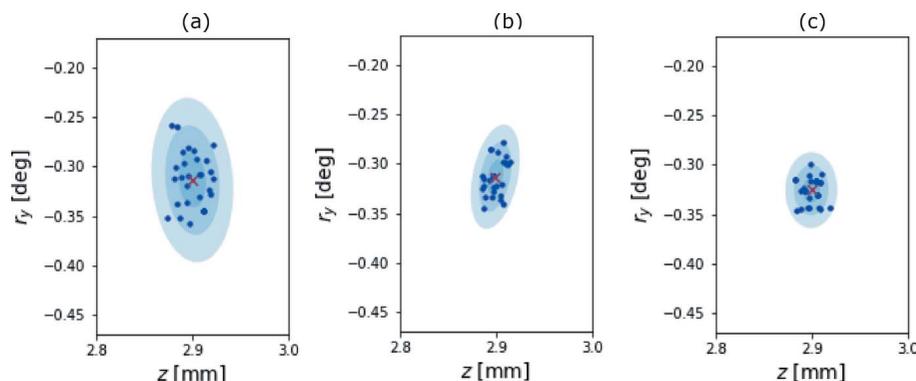
We performed three Monte Carlo experiments, each fixing the value for the noise parameter  $\zeta$  to be  $5.0 \times 10^{-n}$ , for  $n = 2, 3, \text{ and } 4$ . The results are presented as scatter plots of the final ‘best’ alignment position in Fig. 7. We observe that, for noise levels comparable with those seen at APS ( $\zeta = 5.0 \times 10^{-3}$ ), the spatial distribution of alignment results of the numerical Monte Carlo simulation qualitatively agree with those seen at the synchrotron Monte Carlo. As expected, we also see that the corresponding confidence regions appear to grow proportionally with  $\zeta$ .

### 5. Discussion and outlooks

Further exploration of this technique will likely lead to improvements in speed, accuracy, and precision. In particular, we look forward to identifying reliable choices of the parameters required to initialize Algorithm A.1. For example, in



**Figure 7** Depicted are the spatial results of three Monte Carlo simulations, in two of the four total dimensions. We varied the additive noise parameter  $\zeta$  between each sampling. Panel (a) depicts the case where  $\zeta = 5.0 \times 10^{-4}$ , (b) shows the case where  $\zeta = 5.0 \times 10^{-3}$ , and (c) shows the  $\zeta = 5.0 \times 10^{-2}$  case. All results are depicted as blue scatter points above three blue confidence regions. The mean result is presented as a red cross.



**Figure 8** Depicted are the spatial results of three Monte Carlo simulations, in two of the four total dimensions. We included the simulated beam jitter, and fix the additive noise parameter  $\zeta = 5.0 \times 10^{-3}$ , varying the reflection parameter within Algorithm A.1. Panel (a) shows our default choice,  $\alpha = 1.0$ , panels (b) and (c) decrease the parameter to  $\alpha = 0.75$  and  $0.5$ , respectively. The results are depicted as blue scatter points above three confidence regions, and a red cross depicting the mean result.

the settings established for the simulated study in Section 4.3, reducing the reflection parameter to  $\alpha = 0.75$ , or further to  $\alpha = 0.5$ , in Algorithm A.1 results in substantial improvements to the spatial statistics. These results are seen in Fig. 8. Whether or not these changes improve performance on a beamline is an open question.

Second, since an attractive application of this technique is a reliable unsupervised re-alignment procedure, an investigation into automated stopping conditions is required. We anticipate any such condition to be determined by three conditions: a prescribed maximum tally of motor motions, a processing of image data from near or far-field camera sensors, and a time-series study of all evaluations of the FOM.

A robust sensitivity study of the input parameters of Algorithm A.1 on the modeled FOM (9) is unlikely to identify optimal parameter settings in real experimental configurations. However, we suspect that efficient and effective parameter selections can be identified with moderate calibration time.

## 6. Conclusions

In this paper we presented a technique to automate the alignment of CRLs. The technique was implemented and tested at the Advanced Photon Source at Argonne National Laboratory. The complete algorithm was presented, along with a study of its efficacy in synchrotron applications using nominal settings. Additionally, we verify and reproduce the experimental tests in a simulated repeatability study. Lastly, we propose a number of improvements which can be verified with more time at an X-ray light source.

### APPENDIX A

Full details of the algorithms mentioned in the paper are discussed herein, and are enumerated in the order they were introduced above. We again note that examples of the following methods have been provided online (Breckling, 2022).

#### A1. Algorithm A.1. Stochastic Nelder–Mead for CRL alignment

Each major step of the SNM method is indexed by  $k$ , and considers a simplex of positions  $\chi_k = \{\mathbf{p}_{1,k}, \dots, \mathbf{p}_{5,k}\}$ . The reflection, expansion, and contraction parameters  $\alpha$ ,  $\gamma$ , and  $\beta$  are required; nominal values are often assigned to be  $\alpha = 1$ ,  $\gamma = 2$ , and  $\beta = 1/2$ . We additionally require the diameters of the local 4D region  $\mathbf{d}$ , and the cooling parameter  $\varepsilon$ .

*Step 0.* Determine a valid initial simplex of five, non co-linear positions; *i.e.* no more than two points within the simplex should fall on the same straight line in  $\Omega$ . Additionally, ensure light passing through the CRL is present on at least one of the positions in the initial simplex  $\chi_0$ , though preferably all positions. Lastly, evaluate all points in  $\chi_0$  using (6).

*Step 1.* Sort the current simplex  $\chi_k$  from ‘best’ to ‘worst’. Denote the best and worst as  $\mathbf{p}_k^{\max}$  and  $\mathbf{p}_k^{\min}$ , respectively.

**If** the stopping criteria is satisfied, **then** exit the procedure.

**Else** let  $\mathbf{p}_k^{2\min}$  denote the second lowest-ranked position, then

**If** the number of points in  $\chi_k$  exceeds 4, **then** remove  $\mathbf{p}_k^{\min}$  from  $\chi_k$ .

**Else** proceed to *Step 2*.

*Step 2.* Compute a centroid  $\mathbf{p}_k^c$  of all remaining points in  $\chi_k$ . Generate a new point  $\mathbf{p}_k^{\text{ref}}$  by reflecting  $\mathbf{p}_k^{\min}$  through  $\mathbf{p}_k^c$  such that

$$\mathbf{p}_k^{\text{ref}} = (1 + \alpha)\mathbf{p}_k^c - \alpha\mathbf{p}_k^{\text{worst}}.$$

Evaluate  $\mathcal{F}(\mathbf{p}_k^{\text{ref}})$ .

**If**  $\mathcal{F}(\mathbf{p}_k^{2\min}) < \mathcal{F}(\mathbf{p}_k^{\text{ref}}) \leq \mathcal{F}(\mathbf{p}_k^{\max})$ , **then** add  $\mathbf{p}_k^{\text{ref}}$  to the simplex; *i.e.*

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{ref}}\},$$

and return to *Step 1*.

**Else-If**  $\mathcal{F}(\mathbf{p}_k^{\text{ref}}) > \mathcal{F}(\mathbf{p}_k^{\max})$ , **then** expand the reflection point as

$$\mathbf{p}_k^{\text{exp}} = \gamma\mathbf{p}_k^{\text{ref}} + (1 - \gamma)\mathbf{p}_k^c.$$

Evaluate  $\mathcal{F}(\mathbf{p}_k^{\text{exp}})$ , then

**If**  $\mathcal{F}(\mathbf{p}_k^{\text{ref}}) < \mathcal{F}(\mathbf{p}_k^{\text{exp}})$ , **then** add  $\mathbf{p}_k^{\text{exp}}$  to the simplex; *i.e.*

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{exp}}\},$$

and return to *Step 1*.

**Else** add  $\mathbf{p}_k^{\text{ref}}$  to the simplex; *i.e.*

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{ref}}\},$$

then return to *Step 1*.

**Else-If**  $\mathcal{F}(\mathbf{p}_k^{\min}) < \mathcal{F}(\mathbf{p}_k^{\text{ref}}) \leq \mathcal{F}(\mathbf{p}_k^{2\min})$ , **then** we compute an outside contraction point, *i.e.*

$$\mathbf{p}_k^{\text{cont}} = \beta\mathbf{p}_k^{\text{ref}} - (1 - \beta)\mathbf{p}_k^c,$$

then evaluate  $\mathcal{F}(\mathbf{p}_k^{\text{cont}})$ .

**If**  $\mathcal{F}(\mathbf{p}_k^{\text{ref}}) < \mathcal{F}(\mathbf{p}_k^{\text{cont}})$ , **then** add  $\mathbf{p}_k^{\text{cont}}$  to the simplex; *i.e.*

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{cont}}\},$$

then return to *Step 1*.

**Else** perform the local Sobol search (Algorithm A.2) to find  $\mathbf{p}_k^{\text{LSS}}$ , and update the simplex

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{LSS}}\},$$

then return to *Step 1*.

**Else**  $\mathcal{F}(\mathbf{p}_k^{\text{ref}}) \leq \mathcal{F}(\mathbf{p}_k^{\min})$ , **then** we compute an internal contraction point, *i.e.*

$$\mathbf{p}_k^{\text{cont}} = \beta\mathbf{p}_k^{\text{ref}} + (1 - \beta)\mathbf{p}_k^c,$$

and evaluate  $\mathcal{F}(\mathbf{p}_k^{\text{cont}})$ .

**If**  $\mathcal{F}(\mathbf{p}_k^{\text{ref}}) < \mathcal{F}(\mathbf{p}_k^{\text{cont}})$ , **then** add  $\mathbf{p}_k^{\text{cont}}$  to the simplex; *i.e.*

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{cont}}\},$$

then return to *Step 1*.

**Else** perform the local Sobol search (Algorithm A.2) to find  $\mathbf{p}_k^{\text{LSS}}$ , and update the simplex

$$\chi_{k+1} \leftarrow \chi_k \cup \{\mathbf{p}_k^{\text{LSS}}\},$$

then return to *Step 1*.

#### A2. Algorithm A.2. Local Sobol search

Given a maximal sample count  $N \in \mathbb{N}$ , the current SNM iteration  $k$ , a given cooling parameter  $\varepsilon$ , the initial box-width estimate  $\mathbf{d}$ , the simplex  $\chi$ , and known evaluations of  $\mathcal{F}(\mathbf{p}_i)$  for all  $\mathbf{p}_i \in \chi$ , we identify a local region  $\hat{\Omega}$  and perform space-filling sample.

*Step 1.* Generate the local search region by creating the box  $\hat{\Omega}$  around the current motor position  $\mathbf{p}$  by identifying the corners (7), and (8). Create an  $N$ -point Sobol sampling of  $\hat{\Omega}$ , denoted as  $\mathcal{S} = \{\mathbf{p}_i^s\}_{i=1}^N$ .

*Step 2.* From the current position of all four stepper motors, compute an approximate shortest path through each untested point of  $\mathcal{S}$  using the nearest-neighbor method.

*Step 3.* Begin evaluating each point within the Sobol sampling sequentially, according to the ordering established in *Step 2*, using the sample-averaged FOM (6)

**If** at any point  $\mathcal{F}(\mathbf{p}_i^s) > \mathcal{F}(\mathbf{p}^{\min})$ , **then** return  $\mathbf{p}_i^s$ .

Else proceed to *Step 4*.

*Step 4.* If no position within  $S$  improves the simplex  $\chi$ , then extend the Sobol sample by  $N$  additional points, and return to *Step 2*.

### A3. Algorithm A.3. Nearest neighbor

Given a list of spatial coordinates  $S$  in  $\Omega \subset \mathbb{R}^4$ :

*Step 0.* Denote the nearest-neighbor solution as the set  $\mathcal{R}$ , and initialize as the empty set  $\emptyset$ . Given the initial position of the motors  $\mathbf{p}_0$ , determine the nearest position within  $S$  to  $\mathbf{p}_0$ , and denote that point as  $\mathbf{p}^*$ .

*Step 1.* Remove  $\mathbf{p}^*$  from the set  $S$ . Let  $m$  denote the number of positions contained in  $\mathcal{R}$ . Enumerate  $\mathbf{p}^*$  as  $\mathbf{p}_{m+1}$ , and add it to  $\mathcal{R}$ .

*Step 2.* If any positions remain in  $S$ , determine the nearest position in  $S$  to  $\mathbf{p}_{m+1}$ . Denote that position as  $\mathbf{p}^*$ , and return to *Step 1*.

### Acknowledgements

This manuscript has been authored in part by Mission Support and Test Services, LLC, under Contract No. DE-NA0003624 with the US Department of Energy, National Nuclear Security Administration (DOE-NNSA), NA-10 Office of Defense Programs, and supported by the Site-Directed Research and Development Program. The US Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published content of this manuscript, or allow others to do so, for United States Government purposes. The US Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>). The views expressed in the article do not necessarily represent the views of the US Department of Energy or the United States Government. DOE/NV/036240-1104. Portions of this work were performed at High Pressure Collaborative Access Team (HPCAT; Sector 16), Advanced Photon Source (APS), Argonne National Laboratory. HPCAT operations are supported by the DOE-NNSA's Office of Experimental Sciences. The Advanced Photon Source is a DOE Office of Science User Facility operated for the DOE Office of Science by Argonne National Laboratory under Contract No. DE-AC02-06CH11357. Part of this work was performed under the auspices of the US Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

### Funding information

Funding for this research was provided by: U.S. Department of Energy, National Nuclear Security Administration (contract No. DE-NA0003624 to Sean Breckling, Marylesa Howard, Arnulfo Gonzales, Ajanae Williams); US Department of Energy, National Nuclear Security Administration (contract

No. DE-AC02-06CH11357 to Paul Chow); US Department of Energy, National Nuclear Security Administration (contract No. DE-AC52-07NA27344).

### References

- Breckling, S. (2022). *Beam Line Optics Tool*, <https://gitlab.osti.gov/brecklsr/blot>.
- Chang, K.-H. (2012). *Eur. J. Oper. Res.* **220**, 684–694.
- Fang, J. & Savransky, D. (2016). *Appl. Opt.* **55**, 5967–5976.
- Fox, B. L. (1986). *ACM Trans. Math. Softw.* **12**, 362–376.
- Kutsal, M., Bernard, P., Berruyer, G., Cook, P. K., Hino, R., Jakobsen, A. C., Ludwig, W., Ormstrup, J., Roth, T., Simons, H., Smets, K., Sierra, J. X., Wade, J., Wattencamps, P., Yildirim, C., Poulsen, H. F. & Detlefs, C. (2019). *Mater. Sci. Eng.* **580**, 012007.
- Lagarias, J., Reeds, J., Wright, M. & Wright, P. (1998). *SIAM J. Optimiz.* **9**, 112–147.
- Lengeler, B., Schroer, C., Tümmeler, J., Benner, B., Richwin, M., Snigirev, A., Snigireva, I. & Drakopoulos, M. (1999a). *J. Synchrotron Rad.* **6**, 1153–1167.
- Li, Z. & Zhan, Y. (2014). *The 4th IEEE International Conference on Information Science and Technology*, 26–28 April 2014, Shenzhen, China, pp. 821–824.
- Nelder, J. A. & Mead, R. (1965). *Comput. J.* **7**, 308–313.
- Ratzloff, J. K., Law, N. M., Corbett, H. T., Fors, O. & del Ser, D. (2020). *J. Astron. Telesc. Instrum. Syst.* **6**, 1–30.
- Schroer, C. G., Kurapova, O., Patommel, J., Boye, P., Feldkamp, J., Lengeler, B., Burghammer, M., Riekkel, C., Vincze, L., van der Hart, A. & Kuchler, M. (2005). *Appl. Phys. Lett.* **87**, 124103.
- Simons, H., Ahl, S. R., Poulsen, H. F. & Detlefs, C. (2017). *J. Synchrotron Rad.* **24**, 392–401.
- Simons, H., King, A., Ludwig, W., Detlefs, C., Pantleon, W., Schmidt, S., Stöhr, F., Snigireva, I., Snigirev, A. & Poulsen, H. F. (2015). *Nat. Commun.* **6**, 6098.
- Sobol, I. (1967). *USSR Comput. Math. Math. Phys.* **7**, 86–112.
- Song, X., Zhang, X., Liu, G., Cheng, X., Li, W., Guan, Y., Liu, Y., Xiong, Y. & Tian, Y. (2011). *Nucl. Instrum. Methods Phys. Res. A*, **659**, 531–536.
- Vaughan, G. B. M., Wright, J. P., Bytchkov, A., Rossat, M., Gleyzolle, H., Snigireva, I. & Snigirev, A. (2011). *J. Synchrotron Rad.* **18**, 125–133.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., Vijaykumar, A., Bardelli, A. P., Rothberg, A., Hilboll, A., Kloeckner, A., Scopatz, A., Lee, A., Rokem, A., Woods, C. N., Fulton, C., Masson, C., Häggström, C., Fitzgerald, C., Nicholson, D. A., Hagen, D. R., Pasechnik, D. V., Olivetti, E., Martin, E., Wieser, E., Silva, F., Lenders, F., Wilhelm, F., Young, G., Price, G. A., Ingold, G., Allen, G. E., Lee, G. R., Audren, H., Probst, I., Dietrich, J. P., Silterra, J., Webber, J. T., Slavič, J., Nothman, J., Buchner, J., Kulick, J., Schönberger, J. L., de Miranda Cardoso, J. V., Reimer, J., Harrington, J., Rodríguez, J. L. C., Nunez-Iglesias, J., Kuczynski, J., Tritz, K., Thoma, M., Newville, M., Kümmerer, M., Bolingbroke, M., Tartre, M., Pak, M., Smith, N. J., Nowaczyk, N., Shebanov, N., Pavlyk, O., Brodtkorb, P. A., Lee, P., McGibbon, R. T., Feldbauer, R., Lewis, S., Tygier, S., Sievert, S., Vigna, S., Peterson, S., More, S., Pudlik, T., Oshima, T., Pingel, T. J., Robitaille, T. P., Spura, T., Jones, T. R., Cera, T., Leslie, T., Zito, T., Krauss, T., Upadhyay, U., Halchenko, Y. O. & Vázquez-Baeza, Y. (2020). *Nat. Methods*, **17**, 261–272.
- Wright, M. (1996). *Numerical Analysis*, pp. 191–208. Harlow: Addison-Wesley.