Check for updates

High Power Laser Science and Engineering, (2026), Vol. 14, e55, 12 pages.
doi:10.1017/hpl.2026.10146

HIGH POWER LASER SCIENCE AND ENGINEERING

RESEARCH ARTICLE

Spatio-temporal characterization of a petawatt laser using deterministic speckle patterns

Jannik Esslinger[1], Sebastian Wiest[1], Nils Weiße[1], Slava Smartsev[2], Stefan Karsch[1], and Andreas Döpp[1]

[1]*Ludwig–Maximilians–Universität München, Garching, Germany*

[2]*Laboratoire d'Optique Appliquée, ENSTA Paris, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France*

(Received 4 February 2026; revised 17 April 2026; accepted 28 April 2026)

Abstract

We present a deterministic method for the spatio-spectral characterization of intense ultrashort laser pulses via far-field speckle analysis. By mapping vacuum propagation to a linear forward model, we replace slow, error-prone iterative algorithms with a direct least-squares minimization. This enables high-speed retrieval of spectrally resolved wavefronts with rigorous uncertainty quantification. Demonstrated at the ATLAS-3000 petawatt facility, our *in situ* architecture utilizes only a pinhole mask for both attenuation and speckle generation. By eliminating complex downstream optics and minimizing systematic errors, this approach provides a robust, real-time diagnostic essential for the optimization of high-intensity lasers.

Keywords: high-power lasers; laser diagnostics

1. Introduction

The development of chirped pulse amplification (CPA) has enabled the construction of laser systems that deliver ultrashort pulses with peak powers reaching the petawatt (PW) level[1,2]. When tightly focused, these pulses can achieve extreme intensities exceeding 10^{23} W/cm^2[3], creating unique opportunities to study laser–matter interactions in the relativistic regime and drive applications such as laser–plasma acceleration[4,5]. Achieving these record intensities is critically dependent on the ability to perfectly focus the laser pulse in both space and time[6]. However, ultrashort pulses are inherently broadband, and dispersive effects accumulated during amplification and propagation can cause different frequency components to possess different wavefronts[6,7].

These effects are known as spatio-temporal couplings (STCs)[6]. Common examples include pulse-front tilt (PFT), often caused by misaligned compressor gratings, and pulse-front curvature (PFC), introduced by chromatic lenses[7,8]. In many cases, STCs are undesirable as they effectively lengthen the pulse duration at the focus, thereby reducing the peak intensity[9]. However, in some advanced applications, STCs are deliberately engineered to control the dynamics of the focal spot, such as creating a 'flying focus' to guide laser–plasma interactions[10,11]. Whether for elimination or control, the precise characterization of STCs is essential for modern high-power laser science.

A variety of techniques have been developed to measure STCs[12], which can be broadly categorized into two classes.

- Some methods rely on spectrally resolved wavefront sensing. Many variations resort to the use of a Shack–Hartmann sensor paired with tunable or narrowband filters, as seen in techniques such as HAMSTER[13] or FALCON[14]. Single-shot capability is achieved using a dispersive element to measure multiple colours at different positions on the sensor[15]. The recently presented RAVEN technique further expands this to vector fields[16].
- Others are based on spatially-resolved spectroscopy, for example using imaging Fourier-transform spectroscopy (FTS). TERMITES[17,18] and INSIGHT[19], which build a hyperspectral data cube by scanning a temporal delay, can be seen as examples for this approach.

Correspondence to: J. Esslinger, Ludwig–Maximilians–Universität München, Am Coulombwall 1, 85748 Garching, Germany. Email: jannik.esslinger@physik.uni-muenchen.de

© The Author(s), 2026. Published by Cambridge University Press in association with Chinese Laser Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (https://creativecommons.org/licenses/by-nc/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

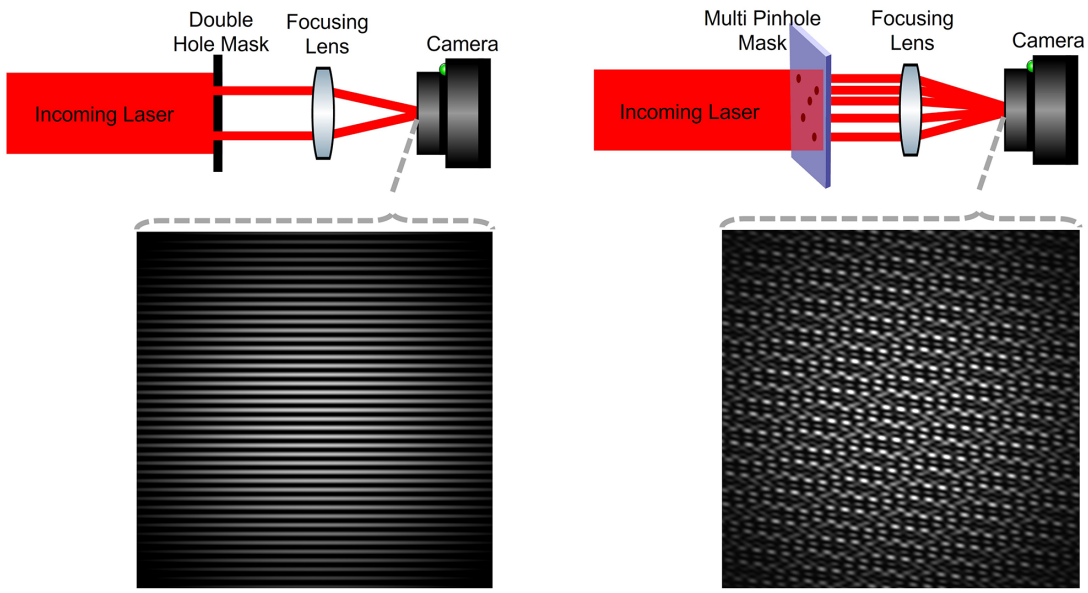


Figure 1. Simplest sketch of the IMPALA setup for a single-colour flat phase laser beam. Also shown are the resulting camera images in the focal plane. It can be seen as an extension to the classical double-slit experiment with holes instead of slits as shown in (a). If there were a phase difference between the beamlets, the fringe pattern would shift by the phase difference. Extending this to five holes with non-regular spacing is shown in (b). The resulting speckle pattern in the focal plane appears chaotic but is in fact the result of a deterministic superposition of individual double pinhole interference patterns.

A somewhat orthogonal approach was recently introduced by Smartsev *et al.*[20]: the IMPALA (Iterative Multispectral Phase Analysis for LAsers) method[20] extends Young's double-slit interferometry (Figure 1(a)) by using a mask with multiple, aperiodically arranged pinholes in the laser's near-field (Figure 1(b)). The interference of the resulting beamlets in the far-field creates a complex speckle pattern that encodes the spectrally resolved wavefront. In the original implementation, this information was retrieved using the iterative Gerchberg–Saxton phase retrieval algorithm[20]. The simple setup allows it to be easily integrated into the beamline without adding any additional dispersing optical setups, making it an *in situ* measurement.

In this paper, we present a new interpretation and a deterministic retrieval algorithm for the setup used for the IMPALA method. Instead of relying on iterative phase retrieval, which might not converge to the correct result, we develop a direct linear forward model based on the theory of multi-beamlet interference in the Fourier domain. This approach allows for a direct, non-iterative reconstruction of the spatio-spectral phase, which is not only computationally faster and more robust but also allows for the propagation of measurement uncertainties. We demonstrate this technique with a single-shot measurement of the ATLAS-3000 PW laser. In our implementation, the pinhole mask serves the dual purpose of creating the required speckle pattern while simultaneously attenuating the full-energy beam, providing a minimal and robust diagnostic. Note that this is an ideal attenuator, as no additional phase is introduced into the transmitted beamlets.

2. Theory

The interference of the focused beamlets can be understood by calculating the field in the focal plane of the focusing optic. As the whole theory can be derived analytically, in parallel, every plane was also simulated in order to verify the formalism as well as to optimize and test the mask design. The general way to calculate the far-field of a focusing optic is given by the Fraunhofer diffraction:

$$E(\vec{r}) = \frac{e^{\left(ikf + \frac{ik}{2f}|\vec{r}|^2\right)}}{i\lambda f} \mathcal{F}\{E_{\text{near-field}}\}\left(\frac{x}{\lambda f}, \frac{y}{\lambda f}\right). \tag{1}$$

Here, $\vec{r}$ is the lateral position vector in the focal plane, f is the focal length, λ is the wavelength with corresponding wave number k and $E_{\text{near-field}}$ is the initial electric near-field. For simplicity, it is assumed that the field is linearly polarized. Explicitly writing the Fourier integral gives the following:

$$E(\vec{r}) = \frac{e^{\left(ikf + \frac{ik}{2f}|\vec{r}|^2\right)}}{i\lambda f} \int_{\text{mask}} E_{\text{near-field}} e^{\left(\frac{ik}{f}\vec{r}\cdot\vec{r}_0\right)} \mathrm{d}\vec{r}_0, \tag{2}$$

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

where $\vec{r}_0$ denotes the lateral position in the near-field. In the case of a mask with infinitely small pinholes, the integral becomes a summation over those holes where the field amplitude is E_m, and the phase is φ_m:

$$E(\vec{r}) = \frac{e^{\left(ikf+\frac{ik}{2f}|\vec{r}|^2\right)}}{i\lambda f}\sum_{m=1}^{N} E_m e^{i\varphi_m} e^{i\vec{k}_m\vec{r}}. \tag{3}$$

Here, $\vec{k}_m = \frac{k}{f}\begin{pmatrix} x_m \\ y_m \end{pmatrix}$ is the wave vector projected onto the focal plane. As a result, the field in the focal plane can be seen as the interference of N plane waves, where N is the number of pinholes.

2.1. Plane wave interference

The individual plane waves are now given by the following:

$$E_m(\vec{r}) = E_m e^{i\left(\vec{k}_m\cdot\vec{r}+\varphi_m\right)}, \tag{4}$$

where the prefactor of the summation has been included in E_m. The intensity of an arbitrary field is given by $I = \frac{\epsilon_0 c}{2}\epsilon_0 \vec{E}\cdot\vec{E}^*$. For the intensity of the interference field $E_{mn} = E_m + E_n$ we obtain the following:

$$I_{mn} = \frac{\epsilon_0 c}{2}\Big(E_m^2 + E_n^2 + 2E_m E_n \cos\left(\Delta\vec{k}_{nm}\cdot\vec{r} + \Delta\varphi_{nm}\right)\Big). \tag{5}$$

In Fourier space this expression results in three distinguished peaks (δ-functions):

$$\mathcal{F}\{I_{mn}\} = \frac{\epsilon_0 c}{2}\Big(\left(E_m^2+E_n^2\right)\delta\left(\vec{\xi}\right) + E_m E_n e^{i\Delta\varphi_{nm}}\delta\left(\vec{\xi}-\Delta\vec{k}_{nm}\right) + E_m E_n e^{-i\Delta\varphi_{nm}}\delta\left(\vec{\xi}+\Delta\vec{k}_{nm}\right)\Big). \tag{6}$$

The latter two terms are the interference terms, which are complex due to the phase shift $\Delta\varphi_{nm} = \varphi_n - \varphi_m$ between the waves. Furthermore, the position of these peaks is solely determined by the difference of the wave vectors $\Delta\vec{k}_{nm} = \vec{k}_n - \vec{k}_m$. As the longitudinal ($z$) component is identical for both, this is simply the distance vector between the pinhole positions in wave numbers. If we denote the positions of the pinholes as (x_m, y_m) and (x_n, y_n), then the vector $\Delta\vec{k}_{nm}$ can be written as follows:

$$\Delta\vec{k}_{nm} = \frac{2\pi}{\lambda f}\begin{pmatrix} x_n - x_m \\ y_n - y_m \end{pmatrix} = \frac{2\pi}{\lambda f}\Delta\vec{x}_{nm}, \tag{7}$$

where λ is the wavelength of the light used in the experiment and f is the focal length of the focusing optic.

We can easily extend this result to multi-beam interference, where the intensity is given as the sum of all combinations of interferences. When considering N beams, the total electric field will be $\vec{E} = \sum_{m=1}^{N} E_m e^{i\left(\omega t - \vec{k}_m\cdot\vec{r}+\varphi_m\right)}$. This results in an intensity of

$$I = \frac{\epsilon_0 c}{2}\sum_{m=1}^{N}\sum_{n=1}^{N} E_m E_n e^{i\Delta\varphi_{nm}} e^{i\Delta\vec{k}_{nm}\cdot\vec{r}}. \tag{8}$$

In Fourier space we can now see the pattern described above with a little modification:

$$\mathcal{F}\{I\}\left(\vec{\xi}\right) = \frac{\epsilon_0 c}{2}\sum_{m}^{N}\sum_{n}^{N}\left(E_m E_n e^{i\Delta\varphi_{nm}}\delta\left(\vec{\xi}-\Delta\vec{k}_{nm}\right)\right). \tag{9}$$

For $n = m$ we again identify the real part E_n^2, because $\vec{k}_{nn} = 0 = \Delta\varphi_{nn}$. For switching n with m we can verify with $\vec{k}_{nm} = -\vec{k}_{mn}$ as well as $\Delta\varphi_{nm} = -\Delta\varphi_{mn}$ that the form of our two complex terms is again given.

Assuming the problem is spatially sparse and can be sufficiently separated into parts that do not overlap, we can isolate the terms, that is

$$\mathcal{F}\{I\}\left(\Delta\vec{k}_{nm}\right) = E_m E_n e^{i\Delta\varphi_{nm}}. \tag{10}$$

From here we can take the argument and log-absolute values to obtain the phase difference and the field amplitude:

$$\arg\left(\mathcal{F}\{I\}\left(\Delta\vec{k}_{nm}\right)\right) = \Delta\varphi_{nm} = \varphi_n - \varphi_m, \tag{11}$$

$$\log\left(\operatorname{abs}\left(\mathcal{F}\{I\}\left(\Delta\vec{k}_{nm}\right)\right)\right) = \log\left(E_m E_n\right) = \log\left(E_m\right) + \log\left(E_n\right). \tag{12}$$

This result implies that we can evaluate the Fourier transform of the interference pattern and retrieve a signal that is directly related to the phase and amplitude of the interfering plane waves.

2.2. Modelling the imaging system

In practice, the interfering waves will have an intensity distribution that is determined by the pinhole shape and the focusing optics. For circular pinholes, the far-field intensity pattern takes the form of an Airy pattern, whose Fourier transformation is again circular. However, we do not deal with perfect imaging systems and the Fourier space expression above has thus to be modified. This can be achieved

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

by convolving the previous equations with the shape of the pinholes, the point spread function (PSF) of the optical system and the PSF of the detector range. Most generally, we can use a function $f(\xi)$ instead of $\delta(\xi)$ to take this form factor into account. The important consequence of the form factor is that the clear spatial separation of features that we relied on so far disappears, at least to some extent. Consider two interference points $\Delta\vec{k}_{nm}$ and $\Delta\vec{k}_{ij}$ whose distance in Fourier space $\| \Delta\vec{k}_{nm} - \Delta\vec{k}_{ij} \|$ is within the extent of the form factor. In this case,

$$\mathcal{F}\{I\}\left(\Delta\vec{k}_{nm}\right) = f(0)E_m E_n e^{i\Delta\varphi_{nm}} + f\left(\| \Delta\vec{k}_{nm} - \Delta\vec{k}_{ij} \|\right) E_i E_j e^{i\Delta\varphi_{ij}} \quad (13)$$

contains contributions from both interferences that add up. Now we cannot perform the same simple separations to retrieve phase and amplitude. One solution to this problem is to only consider points in the retrieval that are clearly separated and discard those that overlap. One can also try reducing the extent of $f(\xi)$, for example by using smaller pinholes in the mask or using a deconvolution with the known form factor. With the simplest approximation, we can think of $f(\xi)$ being a Gaussian distribution centred around $\vec{k}_{nm}$. The standard deviation is determined by the properties of the optical setup: $\sigma^2 \approx \frac{A^2}{(0.84\lambda f)^2} + \frac{1}{D^2}$, where A is the pinhole size, D is the size of the imaging detector and 0.84 comes from approximating the Airy pattern with a Gaussian distribution[21]. With this, it is already clear that reducing the pinhole size will be effective only up to a limit of $A \approx \frac{0.84\lambda f}{D}$. From then on, the positive effect of reducing the pinhole size is on a smaller scale. In addition, one would also have to consider the non-perfect imaging in the form factor, for example due to the finite size of any focusing optic.

2.3. Broadband sources

So far, only the case of one wavelength has been considered. Ultrashort lasers have a broadband spectrum, meaning that $\vec{\Delta}k_{nm}$ as well as $E_{m,n}$ and $\Delta\varphi_{nm}$ depends on λ. Hence, one must integrate over it:

$$\mathcal{F}\{I\} = \sum_{n=1}^{N}\sum_{m=1}^{N}\int E_n(\lambda) E_m(\lambda) e^{i\Delta\varphi_{nm}(\lambda)} \times f\left(\vec{\xi} - \vec{\Delta}k_{nm}(\lambda)\right) \mathrm{d}\lambda. \quad (14)$$

The main consequence is that for broadband sources, streaks are observed, rather than single diffracted points. This acts like the dispersion by a prism or grating and therefore allows direct access to the spectrally resolved phase differences and amplitude products. However, a superposition of streaks at different frequencies is now also possible, which must be considered regarding the separability of amplitude and phase. While the phase part could be linearized using the assumption of a slowly varying phase, the amplitude cannot be easily separated. Therefore, streak overlap should be excluded as much as possible by the mask design.

3. Retrieval

The theory discussed in the previous section can be used to set up a forward model of the measurement process, which we can then use to retrieve information from experiments.

3.1. Phase retrieval

Retrieving the phase differences is in principle as simple as Equation (11) indicates under the assumption that there is no overlap between streaks from different pinhole pairs, as discussed in the previous section. However, there is one additional condition that must be fulfilled, and that is the intra-streak overlap. As the intensity is measured in the far-field on a finite, pixelated detector, we also only obtained a finite, discrete FT, in practice. Hence, we cannot measure the phase differences depending on the wavelength continuously, but instead have to probe for discrete values/pixels that integrate the field over their size. Hence, it is necessary that $E_n E_m$ and $\Delta\varphi_{mn}$ vary slowly in this range for broadband sources. This condition is usually ensured by the fact that we observe focusable, propagating laser pulses rather than arbitrary electric fields. Note that as the streaks scale with the distance between pinholes, longer streaks also obtain more measurable data points; hence, for those, we can achieve a higher spectral resolution.

3.2. Phase unwrapping

For the following derivations, we assume that we have obtained the phase difference values $\Delta\varphi_{nm}$ for discrete wavelengths and neglect the finite resolution of the sensor. Any phase value has the problem that it is only unique within the interval $[0, 2\pi)$ or equally $[-\pi, \pi)$ and otherwise will be wrapped. For a large beam, this is a plausible case, especially since negative values are also possible, thereby reducing the range of absolute values to π. Hence, one has to make sure that whenever such a phase wrapping occurs, it is detected and corrected for. This can be achieved by using the abundance of phase difference measurements compared to the number of pinholes. Given N pinholes, there will be $N \cdot (N-1)/2$ individual pinhole combinations, as these are the triangle numbers. In total, there are $N \cdot (N-1)$ streaks in the Fourier domain due to the indistinguishability of their orientation.

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

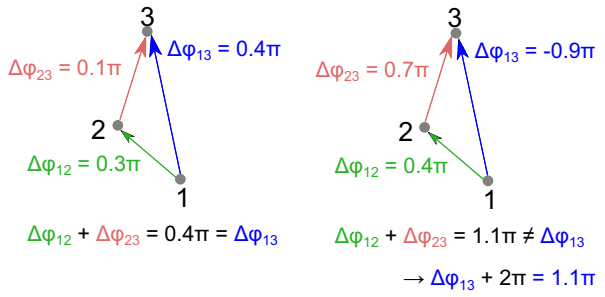


Figure 2. Examples for the phase difference matching test for three pinholes using the phase wrapping of $\Delta\varphi \in [-\pi, \pi)$. In the left-hand example, it checks out, whereas in the right-hand one, it does not, and a correcting factor of $1 \times 2\pi$ has to be taken into account.

As there is no absolute phase as a reference, one has to make the assumption that the pinholes in proximity to each other on the mask do not have an absolute phase difference larger than π (to ensure uniqueness of positive and negative values). From the closer pinholes, we can then deduce if pinholes further apart have wrapped the phase or not by using the additivity of the phase difference values:

$$\Delta\varphi_{nm} = \Delta\varphi_{ni} + \Delta\varphi_{im}. \tag{15}$$

This is easier to understand using an example: consider three pinholes as shown in Figure 2. Pinholes 1 and 2 are the closest, and 2 and 3 are the second closest, whilst 1 and 3 are the furthest apart. In theory, the obtained phase differences of pinholes 1–2 and 2–3 should add up to give the value of 1–3. If this is not the case, one can check which multiple of 2π has to be added to achieve the best match. Of course, for error-free measurements, this is a trivial task and does not provide additional insights; however, in the presence of measurement uncertainties, it will never result in a perfect match and small deviations from a perfect 2π wrapping will occur. If this is taken into account and also near-perfect matches are allowed and correct with a factor of 2π, more measurements can be included in the evaluation to reduce the reconstruction uncertainty. Extending this check to all pinhole combinations yields a net of connected measurements that allows for deducing the unwrapped phase differences for all pairwise pinhole interferences.

3.3. Reconstruction of the phase map from phase differences

The phase values φ_n can easily be calculated, as they are related to the differences by $\Delta\varphi_{nm} = \varphi_n - \varphi_m$. Without measurement uncertainty and other sources of errors, the approach in the previous section should allow for the direct deduction of the individual phase values for each pinhole. As there are more phase differences measured than the number of pinholes, the system is overdetermined, and most of them are only required to verify the final result. However, in the presence of measurement noise, these additional measurements allow for a least-squares solution of the problem: combining all $\Delta\varphi_{nm}$ and all φ_n into two vectors $\Delta\varphi$ and φ, respectively, the relation between the phase values and phase differences can be expressed as a linear algebra problem:

$$\Delta\varphi = \mathbf{R}\varphi. \tag{16}$$

Each row of $\mathbf{R}$ represents the subtraction of a pair of phase values. For an exemplary case of three pinholes, the matrix $\mathbf{R}$ is defined as follows:

$$\mathbf{R} = \begin{pmatrix} +1 & -1 & 0 \\ -1 & +1 & 0 \\ -1 & 0 & +1 \\ +1 & 0 & -1 \\ 0 & +1 & -1 \\ 0 & -1 & +1 \end{pmatrix}. \tag{17}$$

As long as the overlap is sparse, this equation system is overdetermined, and a matrix inversion is not possible. Yet, with the pseudo-inverse $\mathbf{R}^+$ the solution having a minimal root mean squared error (RMSE) can be found:

$$\varphi = \mathbf{R}^+\Delta\varphi. \tag{18}$$

This way of retrieving the phase values in the near-field has a particularly neat side effect: one can leave out any number of phase differences as long as there are still more than the number of phase values to be reconstructed (the matrix stays overdetermined). Of course, the reconstruction accuracy decreases when doing this. On the other hand, it has the advantage that one could fit more pinholes on the mask, such that it becomes unavoidable that some streaks in the FT of the intensity in the focal plane overlap. As long as this affects only a limited number of streaks, one can just ignore those in the reconstruction, such that with the remaining streaks a full reconstruction is still possible.

3.3.1. Propagation of errors

As each value $\Delta\varphi_{nm}$ is affected by an individual error, it is advantageous for the reconstruction accuracy to extract as many phase differences as possible. However, the linear equation systems allow us to estimate the actual uncertainty on the reconstructed phases. The uncertainty on the measured phase values can be estimated by taking multiple images of the intensity in the focal plane in a short time span, in order to assume that the laser beam parameters are constant for all images. One can then estimate the uncertainty accordingly by comparing the resulting phase differences for each image. The result is a diagonal covariance matrix $\Sigma^2_{\Delta\varphi}$. The corresponding covariance matrix for the phase map Σ^2_{φ}

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

can then be estimated via the following:

$$\Sigma_{\varphi}^{2} = \mathbf{R}^{+}\Sigma_{\Delta\varphi}^{2}\mathbf{R}^{+\mathrm{T}}. \tag{19}$$

3.3.2. Modal fitting

In principle, one can stop here; however one can also fit a modal base to the reconstructed phases, which allows for interpreting the result in a physically easier way as well as estimating the phase map for the full beam. If fewer modes are reconstructed than measurements, this further reduces the influence of noise. A particularly suitable base for fitting the wavefront and STCs combined is the Zernike–Taylor base as introduced in Ref. [14] and demonstrated in Refs. [16,22]. This set of basis functions combines the Zernike polynomials for fitting the wavefront component with the Taylor expansion for the spectral phase component. The modes are defined as follows:

$$\varphi(x,y,\omega) = \sum_{m,n,i} a_{m,n}^{i}(\omega-\omega_0)^{i} Z_{n}^{m}(x,y). \tag{20}$$

Here, $a_{m,n}^{i}$ are the Zernike–Taylor coefficients to be reconstructed and Z_{n}^{m} are the Zernike polynomials. In order to obtain the $a_{m,n}^{i}$, another matrix can be set up that links these coefficients to the phase values as shown for phase gradients in Ref. [14]. A pseudo-inverse of the resulting forward matrix can be obtained and then easily integrated in the introduced matrix formalism.

3.4. Reconstruction of the amplitude map

Similar to Equation (16), the matrix formalism for the amplitude is given by the following:

$$\log\left(\mathbf{I}_{\text{far-field, FT}}\right) = \mathbf{T}\log\left(\mathbf{E}_{\text{near-field}}\right). \tag{21}$$

Here, each row of $\mathbf{T}$ is the pairwise addition of the logarithm of field amplitudes $\mathbf{E}_{\text{near-field}}$ in order to obtain the intensity in the Fourier domain of the far-field $\mathbf{I}_{\text{far-field, FT}}$. Overall, this is more straightforward than the phase, as no wrapping may occur for the amplitude. Again, for three pinholes $\mathbf{T}$ would be defined as follows:

$$\mathbf{T} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}. \tag{22}$$

4. Experiments

4.1. General considerations

From the previous discussion we have seen that one of the key parameters determining the resolution of the device is the convolution of the signal due to the form factor. It would seem that a smaller pinhole is always better, but in practice we have to be able to resolve it.

The inference pattern at the focus position f emerging from at least two beamlets separated by a distance d has a periodicity of $\lambda f/d$ and it should be sampled with half that period. For $f/d = 50$ and $\lambda = 800$ nm this means the features of the speckle pattern will occur at a periodicity of 40 μm and, thus, sampling should occur with 20 μm resolution or higher.

Meanwhile, the angle resolution, and thus the ability to resolve the pinholes themselves, is determined by the size of the sensor. The far-field distribution introduced by the pinholes is chromatic and the smallest resolvable angle $\theta_{\min}$ can be estimated by the Rayleigh criterion $\theta_{\min} \approx 1.22\lambda/A$, where A is the aperture size. To resolve a pinhole with the size of the sensor D at a focal length f, we require that $\theta_{\min} \leq D/f$. Substituting the expression for the diffraction limit and rearranging for the largest pinhole size given the wavelength, focal length and sensor size, we get $A \geq 1.22\lambda f/D$. For 800 nm wavelength with $f = 10$ m and a sensor size of $D = 10$ mm, we get $A_{\min} \approx 1$ mm.

So we get a minimum resolution $\Delta x \leq \frac{1}{2}\lambda f/d$ and a minimum sensor size $D \geq 1.22\lambda f/A$. This translates into a pixel number $n_{\text{px}} \geq 2.5d/A$. With a fixed beam size, as well as pixel size and number on the camera, this relationship allows us to determine the ideal magnification (or de-magnification).

4.2. Mask design

An important part of the experimental design is the mask design and how the number and arrangement of pinholes are chosen. As the speckle pattern in the far-field is deterministic, we can directly influence how the far-field and its Fourier transform look. For obvious reasons, a mask containing a regular grid of pinholes or otherwise recurring pattern does not work as streaks in the Fourier domain overlap and a direct, linear retrieval is not possible. The same holds for a completely random arrangement of pinholes. For a small number of pinholes, the chances of a clear separation in the Fourier domain are quite high, but on the other hand one might not cover enough of the original beam. For a larger number, the opposite is the case, and no trade-off might be sufficient.

Therefore, one wants to set certain criteria for the mask to optimize the pinhole pattern for the given number

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

of pinholes. This can be achieved by quantifying the requirements and weighing them appropriately for the setup in order to build a merit function for the mask. If one uses a true random or semi-ordered mask as a starting point, one can move the positions in order to maximize the merit function. In this way, we transform an initial pattern into a pseudo-random pattern. Requirements that need to be considered are as follows.

(1) Rather trivial is that for any mask in which the pinholes are within the diameter of the beam.

(2) The pinholes should be distributed near-uniformly on the mask in order to create a phase map with a suitable distribution of measured phase values. As an extreme example, it would not be sensible to measure only phase values in the upper half of the beam and none in the lower half.

(3) A minimum distance of neighbouring pinholes ensures that the redundancy of information is minimized. Beamlets from neighbouring pinholes typically do not have a large difference in phase and amplitude.

(4) For obvious reasons, overlap in the Fourier domain should be minimized. This ensures that the maximum information can be utilized.

(5) An even distribution of streaks in the Fourier domain ensures that features in any orientation can be observed. For a large number of pinholes that are evenly distributed in the near-field, this criterion is fulfilled almost automatically.

(6) If the number of pinholes is large enough that some overlap in the Fourier domain is unavoidable, one wants to make sure that the lost information affects the reconstruction quality as little as possible. Therefore, one could take the condition number of the inverse matrix in Equation (18) into account.

In addition, it is beneficial to measure as many phase and amplitude values as possible.

4.3. Calibration

For a predesigned, pseudo-random pinhole mask, we can therefore not only simulate the image in the far-field but also know at which positions in the Fourier space we need to extract the phase differences and amplitudes. However, in practice, there will be slight differences; for example, if the mask is not placed perfectly perpendicular to the beam path, then the pinhole pattern is effectively skewed, and therefore the streak positions in the Fourier plane are also skewed. Analogue to this would be a rotation of the mask around the optical axis. There are two ways to resolve this. The first one is of a numerical nature, as one could compute the best affine transform between a simulated streak pattern in the Fourier domain and the measured one. However, if the infrastructure allows for it, a continuous-wave (CW) laser beam that follows the same beam path as the high-power laser pulse to be measured can and should be used as a reference. For this, in the resulting Fourier transform of the far-field intensity, there are only spots instead of streaks, for which it is rather easy to determine their position, for example by using a peak finder or a centre of mass detection. The resulting positions can then be matched to pinhole pairs from the pre-calculated pattern, for example by a linear sum assignment algorithm. From there, we can directly calculate the expected streak pattern, since the position of the other wavelengths directly scales with the ratio of wavelengths $k = \frac{\lambda_{cw}}{\lambda} k_{cw}$. This works especially well in this scenario, where there is no additional imaging involved that may introduce aberrations.

4.4. Experimental setup

The ATLAS-3000 laser system is a titanium:sapphire (Ti:Sa) laser and is capable of delivering a maximum pulse energy of 90 J before compression. After compression, the laser pulse duration is 27 fs, placing its peak power in the PW range. After the final expander, the beam diameter measures 27 cm.

For particle acceleration experiments, the system can provide the laser pulse at full energy or, for diagnostic purposes, at much lower energy levels by using a reflective attenuator positioned after the final amplifiers and before the compressor and beam expander. However, the compressor not only scales down the overall laser pulse energy but also alters its spatio-temporal shape. Therefore, the use of the pseudo-random mask is a more elegant solution to attenuate the beam for diagnostics, as the transmitted laser beamlets do not pass any optical element as long as the hole size is large compared to the hole length and the laser wavelength. In addition, in this way it is also shot at full power on the compressor gratings, allowing for the measurement of dynamic effects by compressor heating as well[23]. In the conducted experiments, the laser operated at 19.5 J prior to compression.

We designed a mask containing 18 holes for the demonstration of our technique. For the optimization of the pseudo-random pinhole pattern, we used the mentioned possible criteria (1)–(5) from Section 4.2. For the chosen amount, full separation of streaks in the Fourier domain is still possible, and hence criterion (6) is not necessary. The beam radius for the maximum size of the pinhole distribution is 28 cm and the size for the pinholes was chosen to be 0.85 mm. The final mask design is shown in Figure 3(a) as well as a simulated focal-plane intensity image in Figure 3(b) and the resulting streak pattern in the Fourier domain in

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

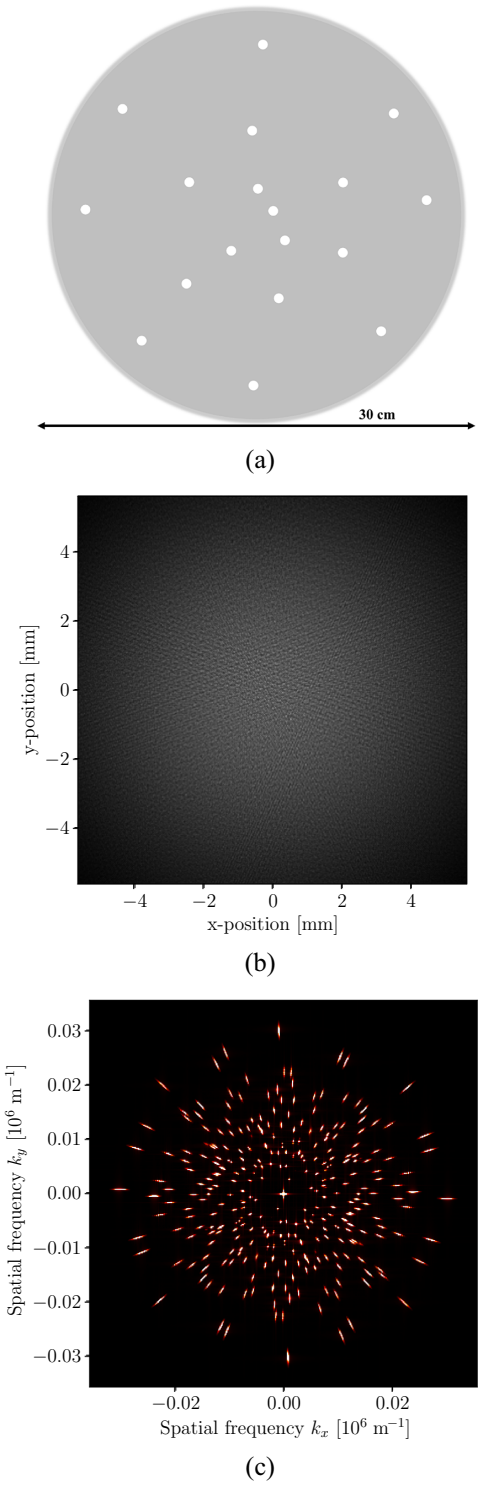


Figure 3. (a) Front-view of the final mask design with 18 pinholes. The pinhole size is significantly exaggerated for visibility. (b) Simulated image in the far-field for the flat phase and amplitude and a Gaussian spectrum. (c) Absolute value of the Fourier-transformed far-field intensity.

Figure 3(c). The necessary nearest neighbour connections for this mask design for monitoring phase wrapping are shown in Figure 4. The mask was manufactured using laser cutting of anodized AlMg3 with a thickness of 0.5 mm.

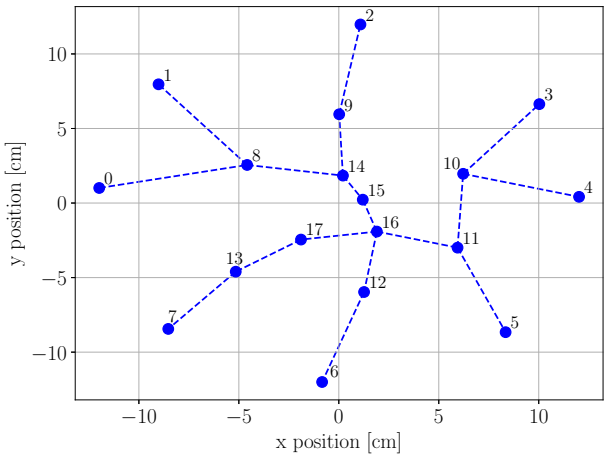


Figure 4. The phase connections used for checking for phase wrapping.

This material was chosen as its reflection is diffuse, which is beneficial for the experimental operation. As only a small part of the laser energy is transmitted, a significant amount of energy must be deposited. An absorbing material would more likely degrade over time, especially when using a reasonable thickness. However, we cannot reflect it directly back as the reflected pulse could go back through the amplification beamline and damage the laser frontend. In addition, at ATLAS-3000, there is a fast Pockels cell installed to block these back-reflections as well as a diode that monitors them.

The mask was installed in the Electron and Thompson Test Facility (ETTF) experimental chamber, which is usually used for laser wakefield acceleration of electrons. There, the mask is placed directly after the first mirror in the vacuum chamber and before an $f = 10$ m focusing parabola. It is attached to a motorized, retractable mount that allows operation with and without the mask. The mask moved in the beam path was still at a slight angle compared to the optical axis to further mitigate the risk of direct back-reflections.

We installed an additional mirror on a motorized stage about 70 cm before the focal plane in order to image the far-field intensity outside of the acceleration setup. We placed a U3-3992SE-M Rev. 1.2. camera by iDS in the new focal plane. The camera was also mounted on a motorized stage along the optical axis in order to find the focal plane using the attenuated beam with the mask moved out of the beam path.

4.5. Results

After conducting the measurements, we noticed that the mask diameter was set too large. We assumed a beam size of 27 cm and designed a maximum mask diameter of 28 cm. However, by investigating the part of the mask where a small amount of material was ablated, we deduced that the actual beam diameter is about 24 cm, which has not been

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

investigated before in this experimental setup. Therefore, as a perfect alignment of the mask with the centre of the beam is hardly possible, we deduced that the outer pinholes may not have been illuminated properly. Hence, we decided not to use the streaks from pinholes at this radius in the final reconstruction. In addition, another pinhole was located directly in the shadow of a small lever holding a tiny pick-off mirror used for spectral phase reconstruction. Therefore, this pinhole had to be excluded from the reconstruction as well. The final reconstruction will only include nine pinholes in the central region of the mask. Nonetheless, this does not devalidate showcasing the general principle of the setup and the analysis formalism.

The attenuation of the mask is twofold. Firstly, it blocks a large amount of the laser pulse, given by the squared ratio of the diameter of the beam (d_{Beam}) and the pinholes (d_{Pinholes}), times the number of pinholes (N). This intensity reduction χ_{blocking} is then $N \cdot d^2_{\mathrm{Pinholes}}/d^2_{\mathrm{Beam}}$. In the focal plane, the main size of the area illuminated also scales with the same ratio, leading to the secondary intensity reduction effect. This has to be scaled again by the number of pinholes, which incorporates the possibly constructive interference of the beamlets. This intensity reduction is hence given as well by $\chi_{\mathrm{spreading}} = N \cdot d^2_{\mathrm{Pinholes}}/d^2_{\mathrm{Beam}}$. The total intensity reduction is hence $\chi_{\mathrm{total}} = \chi_{\mathrm{blocking}} \cdot \chi_{\mathrm{spreading}} = N^2 \cdot d^4_{\mathrm{Pinholes}}/d^4_{\mathrm{Beam}}$. In our scenario, this leads to a reduction of about 5×10^8. Together with an ND-8 filter placed in front of the camera, this led to a total reduction of about 5×10^{16}, enough to bring the original intensity of about 5×10^{19} W/cm^2[16] down to below the typical damage threshold of complementary metal–oxide–semiconductor (CMOS) cameras[24]. The attenuation by the mask itself is already strong enough that no nonlinear effects are occurring in the ND filter.

The measured far-field intensity image as well as its Fourier transform is shown in Figure 5. It can be seen that the overlying, zeroth-order Airy pattern fits well on the sensor. Also, in Figure 5(b), some, especially outer, streaks are less bright than the central ones. This indicates that the outer pinholes (for which the maximum distance to other pinholes is greater than for more central pinholes) of the mask have not been illuminated as much as the inner ones, justifying their exclusion from the reconstruction. From there, the reconstructed near-field phase and intensities are shown in Figure 6, proving the feasibility of our presented reconstruction model in this proof-of-principle. Although phase values in the reconstructed phase map seem to rise to larger values at the edge of the beam, this would probably be resolved if all pinholes were included in the reconstruction, showing a flatter phase map. Including more pinholes would also allow one to include higher-order modes in the Zernike–Taylor reconstruction[14], which would allow one to fit higher-order features.

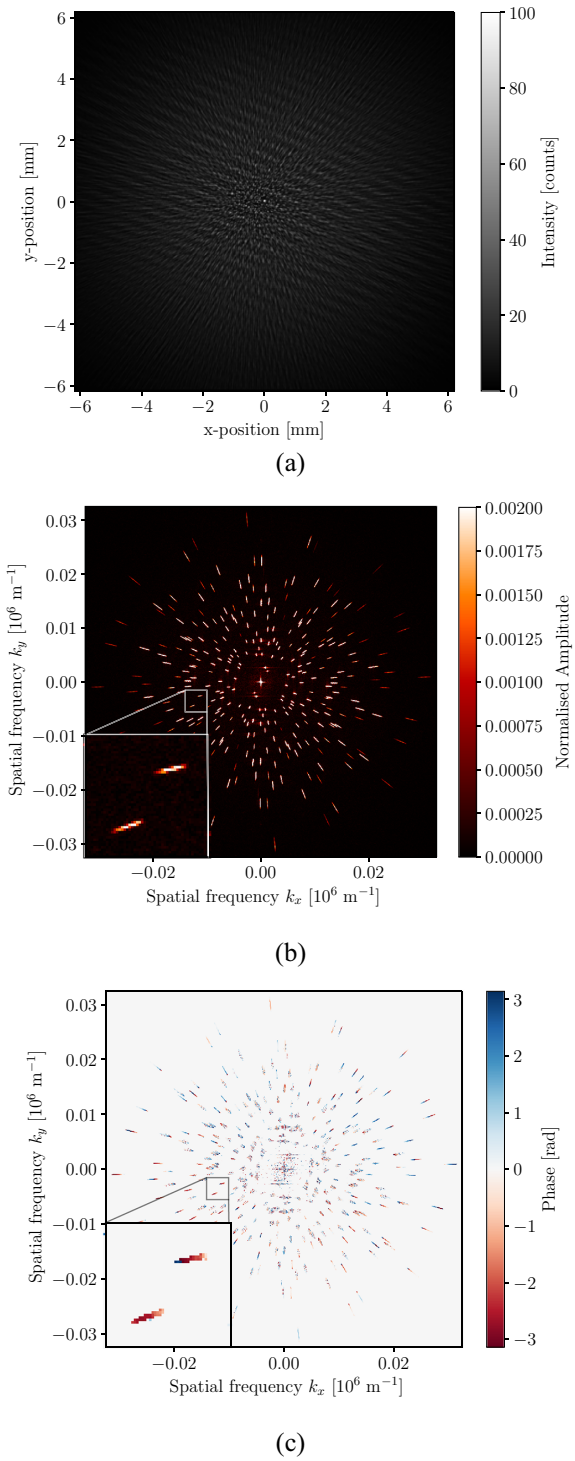
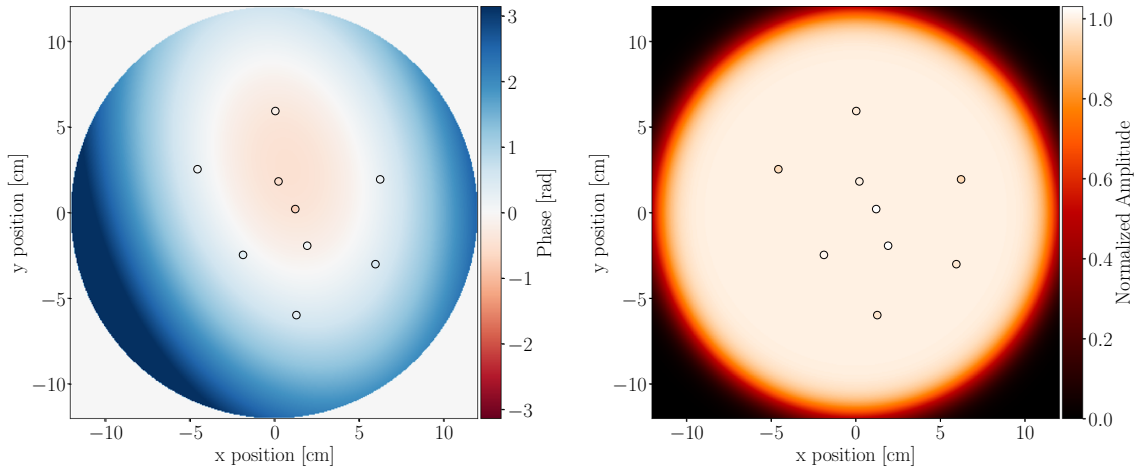
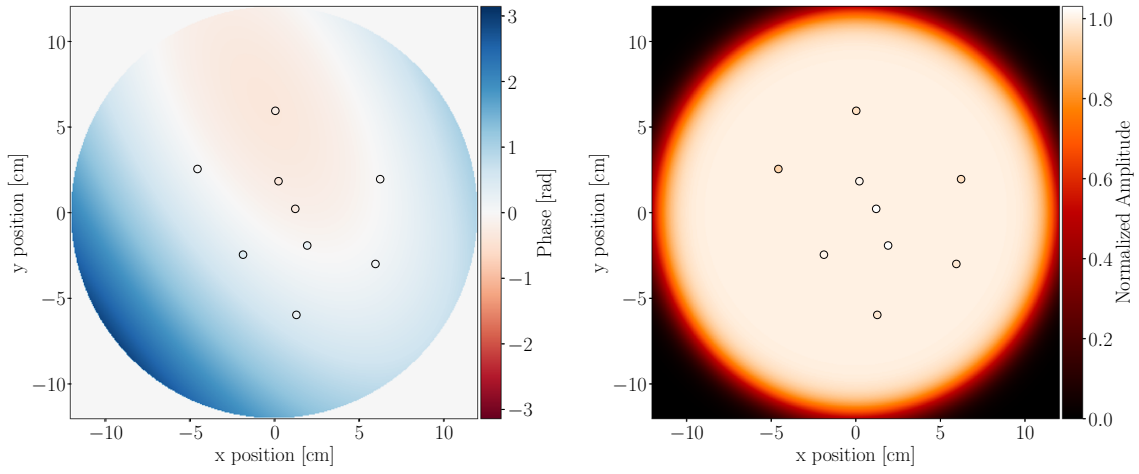


Figure 5. Measured intensity in the focal plane in (a) as well as the amplitude (b) and phase (c) of the Fourier-transformed image. The relative amplitudes in the Fourier domain are rather small due to the prominence of the central peak, which is the direct current (DC) offset. Although this seems to limit the resolution of amplitudes, this dominance is only in the Fourier domain due to the addition of DC peaks from each pinhole without dispersion in one spot. The actual amplitude resolution is mostly limited by the bit depth of the sensor. The phase in (c) was masked such that only physically reasonable values are shown where the amplitude is non-zero.

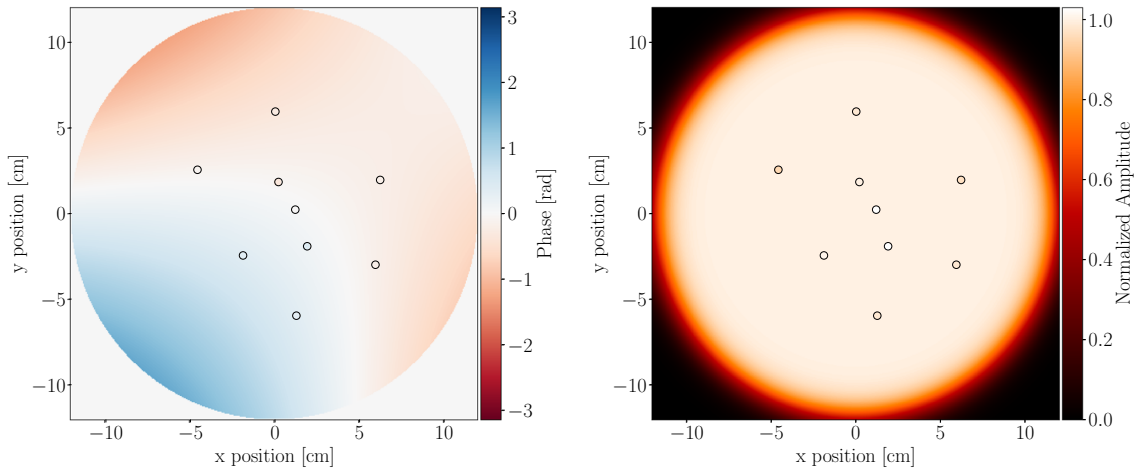
https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press



(a) 785 nm



(b) 800 nm



(c) 815 nm

Figure 6. Reconstructed phase values (left) and relative amplitudes (right) for (a) 785 nm, (b) 800 nm and (c) 815 nm. The reconstructed phase at the pinhole position is shown as well as a full phase map from a low-order Zernike–Taylor modal reconstruction as developed in Ref. [14]. The amplitude values are given relative to their mean value; the flat-top shown is visualizing this.

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

We estimated the uncertainty of each phase difference measurement by taking multiple images in rapid succession and estimated it to be 0.2 rad. Propagating this into the near-field phase space results in uncertainties of about 0.05 rad for each pinhole position. Further using the linear nature of the equations, we can calculate the PFT/angular dispersion to be $0.011 \pm 0.004\ \mu\text{rad nm}^{-1}$, which is in agreement with previous measurements in this facility[14,16,22].

5. Conclusion and outlook

In this work, we have presented a new perspective on the evaluation of the high-power laser far-field generated from a pseudo-random mask originally introduced in Ref. [20]. An analytical forward model enables us to determine where the crucial information, phase differences and combined intensities deduced from the initial near-field are located in the Fourier domain. This is possible as the intensity pattern in the far-field can be directly deduced from the chosen mask design. We then set up linear algebra equations to solve for the original phase values and intensities in the near-field and propagate uncertainties. We showed the option to fit a modal representation of the beam in order to deduce an approximation of the original beam. In addition, complementary methods to the experiment are presented: how to design the mask in order to obtain a viable pseudo-random pinhole pattern; how phase wrapping can be detected and corrected for by using the abundance of measurements; and how to do a final calibration of the mask position and orientation in the experiment. This not only improved the IMPALA technique but also has advantages compared to other established techniques that rely on scanning or iterative reconstruction techniques, such as TERMITES, FALCON and INSIGHT. It now allows single-shot real-time evaluation of laser pulses after the initial calibration.

We successfully tested a mask design using nine pinholes at the ATLAS-3000 laser, where the mask functioned as an attenuator as well through diffuse reflection of the majority of the beam. For the beamlets that passed the mask, no additional dispersing element influenced their properties, allowing for an unaltered reconstruction of the high-power laser pulse. This is especially significant as many diagnostics can only monitor an attenuated beam, which may differ from the original beam.

Future work may include adaptations to the general setup and mask design as well as implementing it in the daily operation routine. For the latter one may use its single-shot capabilities to monitor how the unattenuated beam profile changes over time and compare it to the other studies as shown in Refs. [16,22]. In addition, one can now compare how an attenuator alters the beam profile. In the long term, this setup can be used as a daily calibration of the attenuator in order to correct the results from complementary diagnostics that rely on an attenuated beam. Regarding an extension of the mask design, one can include more pinholes, even until the point of overloading the Fourier domain. This is shown for a potential mask designed using criteria (1)–(5) again, but with 36 pinholes, in Figure 7. Since the system is heavily overdetermined, we can rely on filtering to discard

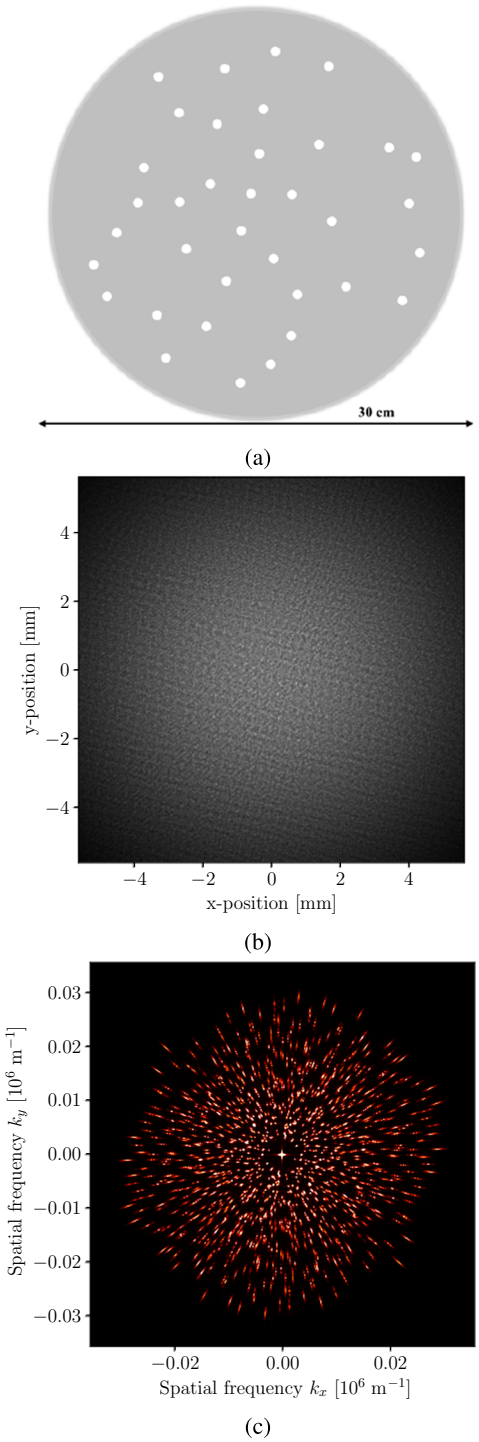


Figure 7. (a) A possible mask design with 36 pinholes. The pinhole size is significantly exaggerated for visibility. (b) Simulated image in the far-field for flat phase and amplitude and a Gaussian spectrum. (c) Absolute value of the Fourier-transformed far-field intensity.

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press

all pinhole pairs that exhibit correlation and only solve the equation system for a discrete set of frequencies and pairwise interferences. Alternatively, one can perform nonlinear retrieval, for example using alternating projections or neural networks based on the forward model, especially for matching the rate of future high-repetition-rate systems. The mask design can also be adapted to avoid outer pinholes not being illuminated properly. This can be achieved by including a safety margin in the maximum diameter of pinhole positions or by including the condition that the pinholes must lie on concentric rings to be excluded automatically if this underillumination is detected. In addition, the experimental setup could be improved by implementing the possibility to move the mask laterally to help with the alignment.

In summary, we showed a new and direct interpretation of the deterministic speckle pattern from a pseudo-random pinhole mask and used a dedicated mask for measuring a high-power laser beam without a dedicated attenuator.

Acknowledgements

We would like to thank André Kalouguine (Laboratoire d'Optique Appliquée, ENSTA Paris, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, France) for fruitful discussions, particularly regarding amplitude retrieval. This work was supported by the Independent Junior Research Group 'Characterization and control of high-intensity laser pulses for particle acceleration', DFG Project No. 453619281.

References

1. C. N. Danson, C. Haefner, J. Bromage, T. Butcher, J.-C. F. Chanteloup, E. A. Chowdhury, A. Galvanauskas, L. A. Gizzi, J. Hein, D. I. Hillier, N. W. Hopps, Y. Kato, E. A. Khazanov, R. Kodama, G. Korn, R. Li, Y. Li, J. Limpert, J. Ma, C. H. Nam, D. Neely, D. Papadopoulos, R. R. Penman, L. Qian, J. J. Rocca, A. A. Shaykin, C. W. Siders, C. Spindloe, S. Szatmári, M. G. M. T. Raoul, J. Zhu, P. Zhu, and J. D. Zuegel, High Power Laser Sci. Eng. **7**, e54 (2019).
2. C. Danson, D. Hillier, N. Hopps, and D. Neely, High Power Laser Sci. Eng. **3**, e3 (2015).
3. J. W. Yoon, Y. G. Kim, I. W. Choi, J. H. Sung, H. W. Lee, S. K. Lee, and C. H. Nam, Optica **8**, 630 (2021).
4. G. A. Mourou, T. Tajima, and S. V. Bulanov, Rev. Mod. Phys. **78**, 309 (2006).
5. E. Esarey, C. B. Schroeder, and W. P. Leemans, Rev. Mod. Phys. **81**, 1229 (2009).
6. S. Akturk, X. Gu, P. Bowlan, and R. Trebino, J. Opt. **12**, 093001 (2010).
7. S. W. Jolly, O. Gobert, and F. Quéré, J. Opt. **22**, 103501 (2020).
8. Z. Bor, Opt. Lett. **14**, 119 (1989).
9. C. Bourassin-Bouchet, M. Stephens, S. de Rossi, F. Delmotte, and P. Chavel, Opt. Express **19**, 17357 (2011).
10. A. Sainte-Marie, O. Gobert, and F. Quéré, Optica **4**, 1298 (2017).
11. D. H. Froula, D. Turnbull, A. S. Davies, T. J. Kessler, D. Haberberger, J. P. Palastro, S.-W. Bahk, I. A. Begishev, R. Boni, S. Bucht, J. Katz, and J. L. Shaw, Nat. Photonics **12**, 262 (2018).
12. B. Alonso, A. Döpp, and S. W. Jolly, APL Photonics **9**, 070901 (2024).
13. S. L. Cousin, J. M. Bueno, N. Forget, D. R. Austin, and J. Biegert, Opt. Lett. **37**, 3291 (2012).
14. N. Weisse, J. Esslinger, S. Howard, F. M. Foerster, F. Haberstroh, L. Doyle, P. Norreys, J. Schreiber, S. Karsch, and A. Döpp, Opt. Express **31**, 19733 (2023).
15. P. Gabolde and R. Trebino, Opt. Express **14**, 11460 (2006).
16. S. Howard, J. Esslinger, N. Weiße, J. Schröder, C. Eberle, R. H. W. Wang, S. Karsch, P. Norreys, and A. Döpp, Nat. Photonics **19**, 898 (2025).
17. M. Miranda, M. Kotur, P. Rudawski, C. Guo, A. Harth, A. L'Huillier, and C. L. Arnold, Opt. Lett. **39**, 5142 (2014).
18. G. Pariente, V. Gallet, A. Borot, O. Gobert, and F. Quéré, Nat. Photonics **10**, 547 (2016).
19. A. Borot and F. Quéré, Opt. Express **26**, 26444 (2018).
20. S. Smartsev, A. Liberman, I. Andriyash, A. Cavagna, A. Flacco, C. Giaccaglia, J. Kaur, J. Monzac, S. Tata, A. Vernier, V. Malka, R. Lopez-Martens, and J. Faure, Opt. Lett. **49**, 1900 (2024).
21. B. Zhang, J. Zerubia, and J.-C. Olivo-Marin, Appl. Opt. **46**, 1819 (2007).
22. J. Esslinger, N. Weiße, C. Eberle, J. Schröder, S. Howard, P. Norreys, S. Karsch, and A. Döpp, Proc. Natl. Acad. Sci. **122**, e2510645122 (2025).
23. V. Leroux, T. Eichner, and A. R. Maier, Opt. Express **28**, 8257 (2020).
24. B. Schwarz, G. Ritt, M. Koerber, and B. Eberle, Opt. Eng. **56**, 034108 (2017).

https://doi.org/10.1017/hpl.2026.10146 Published online by Cambridge University Press