

Supplementary Material

1 Derivations

This section develops the fundamental theory of the Moiré meta-spectrometer. We begin by deriving the dispersion law used throughout the main text, demonstrating an affine relationship between the wavenumber and the relative rotation angle of the two Moiré metasurfaces. We then show that, under uniform angular sampling, the corresponding spectral interval $\Delta\lambda$ scales as λ^2 . This scaling arises directly from the linear dispersion in wavenumber and serves as the geometric basis for the content-adaptive scanner.

1.1 Moiré Metalens Principle and Angle-Wavelength Calibration

The operational principle of Moiré meta-optics relies on the mutual rotation of two complementary metasurfaces [1, 2]. We describe the phase profiles of the two metasurfaces in polar coordinates (r, ϕ) as $\Phi_{\text{MO}_1}(r, \phi)$ and $\Phi_{\text{MO}_2}(r, \phi)$, with corresponding complex transmission functions

$$T_{\text{MO}_1}(r, \phi) = \exp(i\Phi_{\text{MO}_1}(r, \phi)), \quad T_{\text{MO}_2}(r, \phi) = \exp(i\Phi_{\text{MO}_2}(r, \phi)). \quad (1)$$

To enable Moiré tuning, the metasurfaces are designed to be phase-complementary, such that $\Phi_{\text{MO}_2}(r, \phi) = -\Phi_{\text{MO}_1}(r, \phi)$. Rotating the second metasurface by a relative angle θ modifies its transmission to $T_{\text{MO}_2}(r, \phi - \theta) = \exp(-i\Phi_{\text{MO}_1}(r, \phi - \theta))$. The joint transmission of the cascaded stack is therefore

$$T_{\text{joint}}(r, \phi; \theta) = T_{\text{MO}_1}(r, \phi) T_{\text{MO}_2}(r, \phi - \theta) = \exp(i[\Phi_{\text{MO}_1}(r, \phi) - \Phi_{\text{MO}_1}(r, \phi - \theta)]). \quad (2)$$

For small relative rotations, the phase difference in Eq. (2) admits a first-order Taylor approximation,

$$\Phi_{\text{MO}_1}(r, \phi) - \Phi_{\text{MO}_1}(r, \phi - \theta) \approx \theta \frac{\partial \Phi_{\text{MO}_1}(r, \phi)}{\partial \phi}, \quad (3)$$

which shows that rotation converts an angular phase gradient into an effective, rotation-controlled wavefront.

To synthesize a tunable lens, we adopt a separable phase design $\Phi_{\text{MO}_1}(r, \phi) = \Phi_r(r) \Phi_\phi(\phi)$ and seek a net quadratic phase that matches a thin lens of focal length

26 f at wavelength λ ,

$$T_{\text{lens}}(r) = \exp\left(-i\frac{\pi r^2}{\lambda f}\right). \quad (4)$$

27 Choosing $\Phi_r(r) = ar^2$ and $\Phi_\phi(\phi) = \phi$ yields $\partial\Phi_{\text{MO}_1}(r, \phi)/\partial\phi = ar^2$. Substituting into
 28 Eq. (3) and equating to the lens phase in Eq. (4) (the overall sign is fixed by the
 29 rotation sense, equivalently by the sign of a) gives the matching condition

$$ar^2\theta = \frac{\pi r^2}{\lambda f}. \quad (5)$$

30 Accounting for a constant rotation stage offset compared to the Moiré lens pair,

$$ar^2(\theta - \theta_0) = \frac{\pi r^2}{\lambda f}, \quad (6)$$

31 which directly yields the calibration relationship

$$\theta = \left(\frac{\pi}{af}\right)\frac{1}{\lambda} + \theta_0. \quad (7)$$

32 Equation (7) predicts a monotonic inverse dependence between rotation angle and
 33 wavelength for fixed f (i.e., $\theta \propto \lambda^{-1}$), with θ_0 capturing the offset between the opti-
 34 cal and mechanical angular references. Experimentally, we calibrate this mapping
 35 using a set of narrow-linewidth laser lines (Supplementary Note 4) to associate each
 36 wavelength with a unique rotation angle read by the encoder, enabling wavelength
 37 assignment and spectral reconstruction from a controlled rotational scan.

38 1.2 Wavelength-squared scaling of the spectral interval

39 The calibrated dispersion law of Eq. (7) is affine in wavenumber $\nu \equiv 1/\lambda$, with slope
 40 $d\theta/d\nu = \pi/(af)$ independent of λ . Equal commanded angular steps $\Delta\theta$ therefore
 41 correspond to equal wavenumber steps,

$$\Delta\nu = \frac{af}{\pi} \Delta\theta, \quad (8)$$

42 and a uniform sweep in θ is a uniform sweep in ν . Converting to a wavelength interval
 43 via $\nu = 1/\lambda$ gives $\Delta\nu = -\Delta\lambda/\lambda^2$, so

$$\Delta\lambda(\lambda) = \frac{af\lambda^2}{\pi} \Delta\theta. \quad (9)$$

44 The λ^2 scaling is therefore a direct geometric consequence of dispersing linearly in
 45 wavenumber while sampling uniformly in angle, and is independent of any property of
 46 the focusing optic beyond the dispersion law itself. A fixed angular step delivers finer
 47 $\Delta\lambda$ at short wavelengths and coarser $\Delta\lambda$ at long wavelengths, naturally matching the
 48 wavelength dependence of typical atomic and molecular line densities.

2 Adaptive Scanning Algorithm

This section details the content-adaptive Faber–Schauder scanner used throughout the main paper. The acquisition proceeds in three stages: a content-agnostic coarse pilot sweep that initializes the interval tree (Phase 1, Section 2.2); a budget-limited, iterative refinement loop that repeatedly bisects intervals containing high levels of detail (Phase 2, Section 2.3); and a final piecewise-linear reconstruction defined on the resulting non-uniform mesh (Phase 3, Section 2.4). The full procedure is summarized in Algorithm S1, and the parameter choices used in this work are listed in Table S1.

2.1 Notation and the Faber–Schauder detail coefficient

Let $f(\theta)$ denote the measured intensity at rotation angle θ , and $\hat{f}(\theta)$ its piecewise-linear reconstruction from a finite sample set. The acquisition spans the angular domain $[\theta_{\min}, \theta_{\max}]$, which Eq. (7) maps bijectively to the wavelength axis $\lambda(\theta)$. For a subinterval $[\theta_L, \theta_R]$ with midpoint $\theta_M = (\theta_L + \theta_R)/2$, the linear interpolant predicts a midpoint intensity $\tilde{f}(\theta_M) = \frac{1}{2}(f_L + f_R)$. After measuring $f_M = f(\theta_M)$, the deviation from this prediction defines the Faber–Schauder detail coefficient

$$d[\theta_L, \theta_R] \triangleq \left| f_M - \frac{1}{2}(f_L + f_R) \right|, \quad (10)$$

the local second-order finite difference of f . Three properties make d a well-founded refine-or-stop criterion. First, it vanishes on affine intervals: if f is linear over $[\theta_L, \theta_R]$ then $d = 0$, so smooth, slowly varying regions yield negligible detail and stop refining early. Second, it scales quadratically with the interval width and tracks the local curvature of f ; for $f \in C^2$, a Taylor expansion about θ_M yields

$$d[\theta_L, \theta_R] = \frac{h^2}{4} \left| f''(\xi_L) + f''(\xi_R) \right| \leq \frac{h^2}{2} \|f''\|_{\infty, [\theta_L, \theta_R]}, \quad (11)$$

where $h = (\theta_R - \theta_L)/2$ is the half-interval width and ξ_L, ξ_R lie in the left and right sub-intervals. The detail is therefore largest near sharp emission lines and decays as h^2 once the interval is small enough for the interpolant to be accurate, concentrating samples where the spectrum is sharp. Third, d is the magnitude of a wavelet coefficient rather than an empirical heuristic: expanding f in the dyadic Faber–Schauder hat-function basis,

$$f(\theta) = f(\theta_{\min}) \varphi_{0,0}(\theta) + f(\theta_{\max}) \varphi_{0,1}(\theta) + \sum_{\ell \geq 0} \sum_k c_{\ell,k} \psi_{\ell,k}(\theta), \quad (12)$$

where $c_{\ell,k} = f(\theta_{\ell,k,M}) - \frac{1}{2}[f(\theta_{\ell,k}) + f(\theta_{\ell,k+1})]$ is the signed form of Eq. (10). The scanner’s decision metric d is thus exactly the absolute Faber–Schauder coefficient of f on the corresponding dyadic interval [3, 4].

2.2 Phase 1: Pilot sweep

Acquisition begins with N_{pilot} equally spaced angular positions across $[\theta_{\min}, \theta_{\max}]$. We use $N_{\text{pilot}} = 5$ throughout this work; this choice is small enough to leave most of the sample budget for the adaptive phase yet large enough that every level-zero interval

82 has a chance of detecting a peak of width $\gtrsim 1\%$ of the sweep range. The five pilot
 83 positions partition the domain into $N_{\text{pilot}} - 1 = 4$ level-zero intervals, which form the
 84 initial refinement frontier $\mathcal{F}^{(0)}$.

85 To minimise unnecessary motor travel during the pilot, the five positions are visited
 86 in greedy nearest-neighbour order starting from the current stage position θ_{curr} :

$$\theta_{\text{next}} = \arg \min_{\theta \in S_k} |\theta - \theta_{\text{prev}}|, \quad (13)$$

87 where S_k is the set of yet-unvisited pilot positions at step k . On a one-dimensional
 88 domain this heuristic is optimal whenever θ_{curr} lies outside the convex hull of S_0 and
 89 incurs at most a $2\times$ overhead otherwise.

90 2.3 Phase 2: Iterative refinement

91 Refinement proceeds in successive rounds $t = 1, \dots, R$. Each round comprises four
 92 steps.

93 **Step 1 (filter).** The current frontier $\mathcal{F}^{(t-1)}$ is filtered to a candidate set $\mathcal{C}^{(t)}$ by
 94 removing nodes that have either reached the maximum tree depth D_{max} or whose
 95 interval width has dropped below the minimum reproducible motor step $\Delta\theta_{\text{min}}$. Nodes
 96 failing either guard are retained on the frontier but are not eligible for splitting in
 97 subsequent rounds.

98 **Step 2 (pre-measurement of midpoints).** For each candidate interval $[\theta_L, \theta_R] \in$
 99 $\mathcal{C}^{(t)}$ whose midpoint has not yet been probed, the spectrometer records a single mea-
 100 surement at $\theta_M = (\theta_L + \theta_R)/2$. If this midpoint (within a small angular tolerance)
 101 coincides with an already sampled position, the value is retrieved from cache instead
 102 of being remeasured. The detail coefficient $d[\theta_L, \theta_R]$ is then evaluated using Eq. (10).
 103 Every candidate must be measured before ranking: ranking intervals only by the end-
 104 point difference $|f_R - f_L|$ would often miss symmetric structures such as emission lines,
 105 whose endpoints can be nearly identical despite a strong peak in between.

106 **Step 3 (intensity-normalised threshold).** The detail threshold for round t is set
 107 relative to the running peak intensity,

$$\tau_{\text{rel}}^{(t)} = \tau \cdot \max_{\theta' \in \mathcal{S}^{(t)}} f(\theta'), \quad (14)$$

108 where $\mathcal{S}^{(t)}$ is the set of all angular positions measured up to and including round t , and
 109 $\tau \in (0, 1)$ is the single dimensionless control parameter of the scanner. Because the
 110 threshold is rescaled by the running peak rather than fixed in absolute camera units,
 111 τ behaves as a fractional detail-floor: $\tau = 10^{-3}$, for example, means “split any interval
 112 whose midpoint prediction error exceeds 0.1% of the largest measured intensity so far.”
 113 The self-normalising form has the additional property that it absorbs source-brightness
 114 variation across experiments without requiring per-source recalibration.

115 **Step 4 (top- K split, with re-queue).** The candidate set $\mathcal{C}^{(t)}$ is sorted by detail
 116 coefficient in descending order. A budgeted top- K policy then walks the sorted list
 117 and:

- 118 (a) declares an interval a *leaf* (and removes it from the frontier) if its detail is below
 119 $\tau_{\text{rel}}^{(t)}$ or if its bisected child width would fall below $\Delta\theta_{\text{min}}$;
- 120 (b) *bisects* the interval if its detail exceeds $\tau_{\text{rel}}^{(t)}$, by acquiring a new sample at the
 121 midpoint (or reusing the pre-measured midpoint from Step 2) and creating two child
 122 intervals $[\theta_L, \theta_M]$ and $[\theta_M, \theta_R]$ at depth $\ell + 1$;
- 123 (c) *re-queues* the interval onto the next round’s frontier $\mathcal{F}^{(t)}$ once the per-round split
 124 count reaches K , so that high-detail intervals that did not fit in the current round’s
 125 budget are reconsidered in the next round.

126 The top- K budget is structurally analogous to top-down adaptive mesh refinement [5]:
 127 it prevents any single round from spending the entire acquisition budget on one feature
 128 and ensures that detail is distributed across multiple peaks before any one peak is
 129 refined to its depth cap.

130 **Termination.** The refinement loop exits when any of four conditions holds: the fron-
 131 tier is empty (every interval has been declared a leaf); the candidate set is empty (all
 132 remaining intervals have hit D_{max} or $\Delta\theta_{\text{min}}$); the round counter reaches the config-
 133 ured cap R ; or the total sample count reaches the budget B . On termination, the leaf
 134 intervals $\mathcal{L}^{(T)} = \{[\theta_i, \theta_{i+1}]\}$ form a non-uniform partition of $[\theta_{\text{min}}, \theta_{\text{max}}]$.

135 The full algorithm is given in Algorithm S1.

136 2.4 Phase 3: Reconstruction

137 After termination, the cached samples $\mathcal{S} = \{(\theta_i, f(\theta_i))\}$ are sorted in ascending order
 138 of angle and converted to wavelength. Within each leaf interval $[\theta_i, \theta_{i+1}]$ the spectrum
 139 is reconstructed by piecewise-linear interpolation,

$$\hat{f}(\theta) = f(\theta_i) + \frac{f(\theta_{i+1}) - f(\theta_i)}{\theta_{i+1} - \theta_i} (\theta - \theta_i), \quad \theta \in [\theta_i, \theta_{i+1}], \quad (15)$$

140 and the result is mapped to the wavelength axis to give $\hat{f}(\lambda)$. The estimator is *anytime*:
 141 a usable reconstruction is available after every round of refinement, and the sequence
 142 $\hat{f}_1, \hat{f}_2, \dots$ converges monotonically as $|\mathcal{S}|$ grows. For functions of bounded second vari-
 143 ation, piecewise-linear interpolation on an adaptively refined dyadic mesh achieves the
 144 optimal N -term approximation rate among order-one schemes [6, 7], which matches
 145 the regularity class of typical spectra: locally smooth across broad continua, locally
 146 rough at narrow emission lines.

147 2.5 Parameter values used in this work

148 Table S1 lists the control parameters used for every measurement in the main text
 149 and supplementary figures. The same parameter set was applied to every light source;
 150 no per-source retuning was performed.

Supplementary Algorithm 1 Content-Adaptive Faber–Schauder Scanner.

Require: domain $[\theta_{\min}, \theta_{\max}]$, relative threshold τ , sample budget B ; guards: max depth $D_{\max}$, min width $\Delta\theta_{\min}$, splits/round K , max rounds R .

Ensure: samples $\mathcal{S} = \{(\lambda_i, f_i)\}$ and piecewise-linear reconstruction $\hat{f}(\lambda)$.

```

1: Pilot: measure  $N_{\text{pilot}}$  equally spaced angles in nearest-neighbour order (Eq. 13),
   store in  $\mathcal{S}$ , and seed the frontier  $\mathcal{F}$  with the consecutive intervals at depth 0.

2: for  $t = 1, \dots, R$  do                                     ▷ refinement rounds
3:   if  $|\mathcal{S}| \geq B$  or  $\mathcal{F} = \emptyset$  then
4:     break
5:   end if
6:    $\mathcal{C} \leftarrow \{I = [\theta_L, \theta_R] \in \mathcal{F} : \text{depth}(I) < D_{\max}, \theta_R - \theta_L \geq \Delta\theta_{\min}\}$ 
7:   for all  $I \in \mathcal{C}$  do                                       ▷ probe midpoints
8:      $m \leftarrow \frac{1}{2}(\theta_L + \theta_R)$ ; measure  $f(m)$  if not cached
9:      $d_I \leftarrow |f(m) - \frac{1}{2}(f(\theta_L) + f(\theta_R))|$            ▷ Faber–Schauder detail
10:  end for
11:   $\tau_{\text{rel}} \leftarrow \tau \cdot \max_{\theta \in \mathcal{S}} f(\theta)$ ;  $\mathcal{F} \leftarrow \emptyset$ ;  $k \leftarrow 0$ 
12:  for all  $I \in \mathcal{C}$  in order of decreasing  $d_I$  do
13:    if  $d_I < \tau_{\text{rel}}$  or child width  $< \Delta\theta_{\min}$  then
14:      continue                                             ▷ interval resolved
15:    else if  $k < K$  then
16:      add children  $[\theta_L, m], [m, \theta_R]$  to  $\mathcal{F}$  at depth +1;  $k \leftarrow k + 1$ 
17:    else
18:      re-queue  $I$  to  $\mathcal{F}$                                      ▷ carried to next round
19:    end if
20:  end for
21: end for

22: Reconstruct: map  $\theta_i \mapsto \lambda_i$  (Eq. 2) and interpolate  $\hat{f}(\lambda)$  piecewise-linearly over
    $\mathcal{S}$  sorted by  $\lambda$ .
23: return  $\mathcal{S}$ ,  $\hat{f}(\lambda)$ 

```

151 The single dimensionless control parameter τ mediates the trade-off between sam-
 152 ple efficiency and reconstruction fidelity. Smaller values of τ allow more sub-threshold
 153 features to be captured and drive more aggressive refinement, yielding higher-fidelity
 154 reconstructions at the expense of additional measurements. Larger values, in contrast,
 155 enforce stronger pruning and earlier termination, resulting in coarser reconstructions
 156 that require fewer samples. Empirically, we observe that $\tau \in [10^{-6}, 10^{-4}]$ spans the
 157 range in which the adaptive scanner surpasses a uniform sweep at matched fidelity for
 158 all tested source classes. The remaining parameters ($D_{\max}, \Delta\theta_{\min}, K, R$) are hardware-
 159 constrained limits rather than fidelity controls: $\Delta\theta_{\min}$ is determined by the minimum
 160 reproducible increment of the ELL14 rotation stage, $D_{\max}$ and R act as safeguards
 161 against pathological inputs, and $K = 2$ guarantees that fine structure is shared across
 162 multiple peaks before any individual peak is refined to its maximum depth.


Supplementary Table 1 Parameters of the content-adaptive Faber–Schauder scanner. These are the fixed configuration used for the adaptive-versus-uniform benchmark of Supplementary Figs. 4–5; the sample budget B is swept over the listed values, and the detail threshold τ is the only fidelity parameter (the remaining entries are hardware-aware guards). The minimum interval width $\Delta\theta_{\min}$ is a software floor on bisection, not the physical stage resolution.

Symbol	Parameter	Setting key	Value
$[\theta_{\min}, \theta_{\max}]$	Angular sweep range	<code>SWEEP_PARAMETERS</code>	$[-10^\circ, +40^\circ]$
N_{pilot}	Pilot sample count	<code>N_COARSE</code>	5
τ	Detail threshold (fractional)	<code>wtp_tau</code>	10^{-4}
$D_{\max}$	Maximum tree depth	<code>wtp_max_depth</code>	30
$\Delta\theta_{\min}$	Minimum interval width (software guard)	<code>wtp_min_width_deg</code>	0.2°
K	Splits per round (top- k)	<code>wtp_top_k</code>	1
R	Maximum refinement rounds	<code>wtp_max_oscillations</code>	100
B	Total sample budget (swept)	<code>SAMPLING_BUDGET</code>	{8, 16, 32, 64, 128, 256}
t_{settle}	Settle time per measurement	<code>settle_seconds</code>	10 ms

3 Adaptive Scan Results

Supplementary Figures 1–3 replay the content-adaptive acquisition on three sources that span the spectral-difficulty range of the main paper: a broadband continuum carrying a single embedded laser line (Fig. 1), a compact fluorescent lamp combining several sharp mercury lines with a phosphor continuum (Fig. 2), and a near-monochromatic 532 nm laser (Fig. 3). In every case the five-sample pilot resolves only the smooth envelope, and each subsequent round spends its budget where the Faber–Schauder detail is largest (on the line cores and continuum shoulders) while leaving the flat regions sparsely sampled. By the final round all sharp features are recovered at a cumulative budget of $N = 128$ samples, the fidelity a uniform sweep reaches only with substantially more measurements. These experiments are the source-resolved manifestation of the main result: adaptive scanning reconstructs mixed line-plus-continuum spectra with roughly $3\times$ fewer measurements on average (up to $7\times$ than uniform sampling at matched fidelity, with a usable spectrum available after every round.

Adaptive vs Uniform Scanning We compare the adaptive scanner to uniform angular sampling across eleven different sources, covering narrow-band emitters (lasers, LEDs, monochromator setpoints) as well as broadband emitters (a CFL and a combined halogen + 532 nm source). As ground truth $G(\lambda)$, we use a densely sampled uniform sweep over $\theta \in [-10^\circ, +40^\circ]$ ($\lambda \in [384, 1220]$ nm), padding angles outside each source’s emission band to its dark level. All evaluations are performed in simulation: the acquisition software is run unmodified on G , with each measurement returning the ground truth at a slightly perturbed angle (jitter $\sigma_\theta = 0.05^\circ$) and corrupted by a calibrated shot-plus-read camera noise model. For uniform sampling, we distribute a total of $N \in \{8, 16, 32, 64, 128, 256\}$ measurements evenly over the angular range; for the adaptive approach, we allocate the same budgets using the scanner described in Supplementary Note 2 (Algorithm S1). Both methods reconstruct the signal on the dense ground-truth grid via piecewise-linear interpolation, without any super-resolution or




Supplementary Figure 1 Round-by-round adaptive reconstruction of a broadband source with an embedded laser line. Each panel shows the piecewise-linear reconstruction $\hat{f}_t(\lambda)$ (colored fill) on the cumulative sample set at the end of round t , with the 0.01° -step ground-truth scan overlaid in black. Panels are labeled by round (Pilot, Round 2–12) and annotate the cumulative sample count N together with the new samples added in that round (in parentheses). Tick marks beneath each axis indicate the angular positions of the new samples (colored) and prior samples (grey). The pilot sweep places $N = 5$ uniform samples across the full band; subsequent rounds preferentially refine the high-curvature region around the laser line while leaving the smooth continuum under-sampled, recovering both features by round 12 with $N = 128$ total measurements.

parametric model fitting. A single set of hyperparameters (τ , top- k , minimum width) is used for all sources, selected by a grid search that optimizes the average performance across sources.

4 Methods


This section presents the experimental realization of the Moiré meta-spectrometer. We outline the optomechanical construction of the cascaded meta-optic pair and its associated translation and rotation stages, the PSF-based alignment protocol used to co-register MO_1 and MO_2 along a shared optical axis, and the high-precision wavelength calibration that relates the rotation angle of the ELL14 stage to a continuous spectral coordinate. Collectively, these procedures define the hardware platform on which the content-adaptive scanner described in Supplementary Note 2 operates.



Supplementary Figure 2 Round-by-round adaptive reconstruction of a compact fluorescent lamp (CFL). Same convention as Supplementary Fig. 1. The CFL emission combines sharp mercury lines near 436, 546, 578, and 611 nm with a broad phosphor continuum spanning the visible band. The pilot sweep ($N = 5$) provides only a coarse continuum estimate; from Round 2 onward, the adaptive sampler concentrates new measurements at the wavelengths of the mercury lines, locking in their positions within the first few rounds and progressively narrowing their reconstructed widths toward the ground-truth FWHM. By the final round, every mercury line is resolved and the phosphor continuum is faithfully recovered with a cumulative budget of $N = 128$ samples.

4.1 Experimental Setup


The experimental setup consists of two meta-optics placed sequentially along a common optical axis. MO_1 is the meta-optic mounted on a compact two-axis translation stage (LM1XY, Thorlabs), which is then joined to the ELL14 motorized rotation stage using external SM1 thread adapters. These adapters provide the required separation between components and avoid mechanical friction during rotation, while also keeping the optical axis mechanically well-defined through SM1-threaded interfaces. The X–Y translation stage provides a direct way to laterally shift MO_1 while the entire assembly still rotates, which is essential for centering the optic relative to the mechanical axis of the ELL14 piezo-motorized rotation stage. A dovetail translation is integrated at the base of the rotation-stage mount to decrease the separation distance between the two meta-optics. This axial degree of freedom allows us to bring MO_1 and MO_2 closer without disturbing the lateral centering, which is important for compact operation and for reducing propagation effects such as Talbot self-imaging over long inter-element distances.



Supplementary Figure 3 Round-by-round adaptive reconstruction of a 532 nm green laser. Same convention as Supplementary Fig. 1. The source is a single near-monochromatic line at 532 nm — the hardest case for blind sampling, as the $N = 5$ pilot straddles it and recovers almost none of its amplitude. From Round 2 onward the adaptive sampler localises the line and spends essentially all new measurements there, restoring the peak within a few rounds and narrowing the reconstructed width toward the ground-truth FWHM until it is fully resolved at the cumulative budget $N = 128$.

MO₂ is the static meta-lens that is mounted only on an X–Y translation stage so that the centroid of MO₂ can coincide with the principal axis passing through the center of MO₁. In other words, MO₁ defines the rotating reference axis, and MO₂ is then translated to be co-axial with that reference. The MO₂ stage also uses a dovetail z -axis translation stage to reduce the distance between the meta-optics and to fine-tune the axial separation during alignment, especially when the PSF shape indicates residual defocus or spacing-induced changes in the effective system response.


The illumination source is placed approximately 40 mm to the left of MO₁ (approximately at a $2f$ distance) and consists of a circular slit of 50 μm diameter coupled to a fiber probe. This configuration provides an approximately point-like input that is well-suited for monitoring the point spread function (PSF) of the combined system. We use a monochrome camera (FLIR Blackfly grayscale) as the detector, and we adjust the camera position and orientation such that (i) the PSF centroid is close to the center pixel and (ii) the dominant lobes are approximately equal in size and intensity. These conditions not only improve visual interpretability but also provide a practical symmetry check that the sensor plane is approximately perpendicular to the optical axis and that gross tilt is minimized.



Supplementary Figure 4 Reconstruction PSNR versus sample budget for adaptive and uniform angular sampling across eleven light sources. Panels show three lasers (407, 533, 637 nm), four LEDs (blue, green, red, white), two monochromator setpoints (520, 620 nm), a CFL, a mixed halogen + 532 nm source, and the across-source average (bottom right). Blue: adaptive sampling; orange: uniform sampling; dashed: 30 dB reference. Adaptive sampling uses one fixed configuration for all sources and budgets (5-sample pilot, greedy top- $k = 1$, detail threshold $\tau = 10^{-4}$, minimum interval width 0.2°). Averaged over all sources, adaptive sampling beats uniform at every budget—by +5.6 dB on average and up to $\approx +11$ dB at mid budgets—with convergence as $N \rightarrow 256$, where uniform becomes dense.

4.2 Alignment


The optical setup is extremely sensitive to lateral misalignment between the centroids of the two meta-optics. In this system, even small decenter between MO_1 and MO_2 can introduce angle-dependent aberrations and apparent motion of the focal spot, which directly degrades repeatability during rotational scanning. To ensure spectral and spatial stability of the focal point during rotation of MO_1 , we use a PSF-driven visual centering procedure designed to satisfy three alignment objectives:



Supplementary Figure 5 Reconstruction PSNR versus acquisition time for adaptive (WTP) and uniform angular sampling. Same eleven sources and panel layout as Fig. 4, here at a fixed budget of $N = 256$ samples: the PSNR of the cumulative reconstruction is plotted against the modelled acquisition time as samples accrue in measurement order. Blue: adaptive sampling; orange: uniform; dashed line: 30 dB. Because adaptive sampling concentrates measurements on spectral features instead of sweeping uniformly, it reaches a given PSNR in substantially less acquisition time ($\approx 6.7\times$ faster to 30 dB on average); the curves converge once uniform sampling has covered the full range.


- **Rotation-axis centering:** the axis of rotation of the motorized rotation stage coincides with the centroid of MO_1 .
- **Co-axiality:** the centroids of MO_1 and MO_2 lie on the same optical axis.
- **Minimal tilt:** the surfaces of the two metalenses are parallel to each other and to the sensor plane, and approximately orthogonal to the line joining the point source and the sensor plane.

Quasi-incoherent PSF as a stable alignment observable. We illuminate the system with a monochromatic collimated laser module and make it quasi-incoherent



Supplementary Figure 6 Optomechanical assembly of the Moiré meta-spectrometer. **a**, Annotated image of the prototype: the rotating meta-optic MO_1 is mounted on a two-axis translation stage (LM1XY, Thorlabs) coupled to a piezo-motorised rotation stage (ELL14) through SM1 thread adapters; the static meta-optic MO_2 is mounted on its own LM1XY downstream, and the focused PSF is recorded by a monochrome scientific camera (BFS-U3-23S3M, FLIR Blackfly). **b**, Reference PSF through MO_1 alone under quasi-incoherent monochromatic illumination, used as the visual reference for the centring procedure of Supplementary Note 4.2. **c**, Schematic of the optical train: a fibre-coupled point-like source illuminates the cascaded MO_1/MO_2 pair through SM1-threaded mounts; the wavelength brought to focus on the sensor plane is selected by the relative rotation angle θ commanded to the ELL14.

by scrambling the phase with a rotating diffuser, which is a sandpapered translucent polymer sheet attached to a small motor. The reason for scrambling is that a purely coherent laser source generates strong speckle and interference artifacts at the detector, which makes the PSF fluctuate in time and biases any centroid-based centering. The rotating diffuser produces a time-varying phase screen that is averaged by the camera exposure, resulting in a stable PSF profile that is reproducible across repeated measurements. Under this illumination, the system PSF becomes an interpretable and repeatable visual reference that we monitor throughout alignment.



Supplementary Figure 7 PSF-driven alignment of the Moiré meta-optics. **a, b**, Single-metalens PSF through MO_1 before (**a**) and after (**b**) centering MO_1 on the ELL14 rotation axis: we track the PSF centroid over a full rotation (KLT tracker), fit a circle to its path, and translate MO_1 (LM1XY) until the fitted radius is below one binned pixel. **c, d**, Cascaded MO_1/MO_2 PSF (RGB sensor) at $\theta = 1^\circ, 3^\circ, 5^\circ, 7^\circ, 10^\circ$ before (**c**) and after (**d**) co-axially aligning MO_2 to MO_1 : misalignment causes angle-dependent focal drift and asymmetric, fan-shaped lobes; alignment yields a stationary centroid and clean, wavelength-selective intensity modulation. **e**, X- θ line-scan of a white-LED source through the aligned system over fine ($\theta \in [0, 14]^\circ$, left) and coarse ($\theta \in [0, 100]^\circ$, right) ranges; the dominant bright track follows the calibrated λ - θ dispersion.

4.3 Centering MO_1 to the rotation axis (ELL14).

The first objective is to align the centroid of MO_1 with the axis of rotation of the motorized rotation stage. A lateral decenter from the rotation axis produces a characteristic and intuitive signature: during smooth rotation, the PSF centroid does not remain stationary, but instead traces a circular trajectory around the true rotation center in the detector plane. Conceptually, MO_1 rotates about a mechanical axis; if the optic center is offset from this axis, the optical response “orbits” around the true rotation center, and the PSF centroid draws an arc (or a full circle over a complete rotation). This trajectory provides a direct, image-domain indicator of decenter, and it is the basis of our centering procedure.

To quantify this trajectory and reduce observer bias, we use a KLT (Kanade-Lucas-Tomasi) feature tracker to track a stable, high-contrast region of the PSF (typically the bright core or a consistent portion of the fan-like PSF) across frames acquired during a rotation sweep. The tracked points form a 2D trajectory in detector coordinates. We then fit a circle to this trajectory and estimate the rotation center (PSF centroid) from the circle fit. We then use the LM1XY translation mount beneath MO_1 to apply small lateral corrections that move MO_1 toward this estimated center. After each adjustment, we repeat the rotation sweep and tracking. The procedure converges when the fitted radius becomes minimal and, visually, the PSF centroid remains nearly stationary over rotation, indicating that the centroid of MO_1 is well registered to the ELL14 rotation axis.

Co-aligning MO_2 to MO_1 and stabilizing the focus during rotation. Once MO_1 is centered, we align MO_2 using its XY translation stage such that the combined system produces a stable focus whose position does not drift during rotation of MO_1 . This step enforces co-axiality between the two meta-optics: MO_1 defines the rotating optical axis, and MO_2 is translated until the PSF symmetry and centroid stability are maximized. The main visual cue is that the MO_1 – MO_2 pair behaves as an effective convex lens whose focusing behavior is controlled by the relative angle between the two meta-optics. Therefore, as MO_1 rotates about its axis, the image of the point source should remain spatially static on the sensor if MO_2 is properly centered and the two surfaces are approximately parallel. At the same time, the PSF intensity is expected to vary smoothly with angle, with a maximum at the relative angle corresponding to the monochromatic illumination wavelength (consistent with the wavelength-selective response of the meta-optic pair). In our experiments, successful alignment is indicated by (i) negligible lateral drift of the PSF centroid over the full rotation sweep, (ii) smooth and repeatable intensity modulation versus angle, and (iii) stable PSF morphology (no abrupt asymmetry changes) across repeated sweeps.

Finally, after co-alignment, we reduce the $MO_1 - MO_2$ separation using the dovetail z -translations on the MO_1 and MO_2 mounts. This step is performed after lateral centering because changing axial spacing can slightly perturb mechanical stresses and alignment. We iteratively reduce the spacing while confirming that the PSF remains stable and centered, thereby achieving a compact configuration with reduced Talbot-length-related sensitivity while preserving repeatability.

4.4 Precision calibration of the Moiré spectrometer

Wavelength calibration of the spectrometer was performed by acquiring the angular emission spectra of six distinct laser configurations (407.08 nm, 533.09 nm, 636.91 nm, 787.47 nm, 849.72 nm, and 974.16 nm). To mitigate the effects of mechanical backlash and hysteresis in the motorized rotational stage, bidirectional spectral acquisitions—consisting of both clockwise and counterclockwise sweeps—were performed for each reference wavelength. The acquired intensity profiles from both scanning directions were averaged to yield a robust, directionally unbiased signal. Because the discrete angular sampling inherently limits the precision of raw peak identification (i.e., using a simple maximum), we determined the exact central angle of each emission line by fitting a Gaussian distribution to the averaged intensity profile surrounding the region of maximum intensity. The resulting highly precise Gaussian-fitted peak centers (θ_{peak}) were mapped against the theoretical wavenumber ($10^7/\lambda$, in cm^{-1}) of each corresponding laser. Linear regression analysis of these six points revealed a highly consistent calibration function, consistent with the linear wavenumber–angle dispersion of Eq. (7): $\nu [\text{cm}^{-1}] = 358.4601 \theta_{peak} + 11761.6953$, with a coefficient of determination (R^2) exceeding 0.999. This experimentally derived linear map was subsequently rigorously inverted to establish the mapping from raw rotational angle to continuous wavelength ($\lambda = 10^7/(358.4601 \times \theta_{peak} + 11761.6953)$), forming the definitive spectral calibration used across all subsequent commercial light source benchmarks and adaptive sampling routines.

References

- [1] Bernet, S. & Ritsch-Marte, M. Adjustable refractive power from diffractive moiré elements. *Appl. Opt.* **47**, 3722 (2008). URL <https://opg.optica.org/abstract.cfm?URI=ao-47-21-3722>.
- [2] Bernet, S., Harm, W. & Ritsch-Marte, M. Demonstration of focus-tunable diffractive moiré-lenses. *Opt. Express* **21**, 6955 (2013). URL <https://opg.optica.org/oe/abstract.cfm?uri=oe-21-6-6955>.
- [3] Douzi, K., Mammass, D., Wittner, T. & Nouboud, F. Faber-schauder wavelet transform, application to edge detection and image characterization. *J. Math. Imaging Vis.* **14**, 91–101 (2001). URL <https://link.springer.com/article/10.1023/A:1011213914008>.
- [4] Mallat, S. *A Wavelet Tour of Signal Processing* Second edn (Academic Press, 1999).
- [5] Berger, M. J. & Oliger, J. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **53**, 484–512 (1984). URL <https://www.sciencedirect.com/science/article/abs/pii/0021999184900731>.
- [6] DeVore, R. A. Nonlinear approximation. *Acta Numer.* **7**, 51–150 (1998). URL <https://www.cambridge.org/core/journals/acta-numerica/article/>

338 [nonlinear-approximation/C8E028C39B8E2F35F277F58FBFFBB504](#).

- 339 [7] Cohen, A., Dahmen, W. & DeVore, R. Adaptive wavelet methods for elliptic
340 operator equations: convergence rates. *Math. Comp.* **70**, 27–75 (2001). URL
341 <https://www.ams.org/journals/mcom/2001-70-233/S0025-5718-00-01252-7/>.