

Compact Low-Cost Hyperspectral Imaging via Angular-to-Spectral Diversity Conversion

Kazuma Fujiwara^{†1}, Takuya Funatomi^{1,2}, Kazuya Kitano¹,
Yuki Fujimura¹, and Yasuhiro Mukaigawa¹

¹ Nara Institute of Science and Technology, Nara, Japan

² Kyoto University, Kyoto, Japan

{fujiwara.kazuma.fj4, kitano.kazuya, fujimura.yuki,
mukaigawa}@is.naist.jp, funatomi.takuya.2c@kyoto-u.ac.jp

Abstract. Snapshot hyperspectral imaging avoids sequential scanning, but systems that jointly achieve stable reconstruction, low cost, and compact optics remain limited. We present a snapshot hyperspectral imaging system based on angular-to-spectral diversity conversion. A tapered kaleidoscope creates replicated views with distinct incidence directions, and a directly attached birefringent filter converts them into view-channel-dependent spectral transmittances, yielding complementary measurements that better condition the inverse problem for more stable single-shot spectral reconstruction. The system preserves a simple pixel-wise linear model for fast non-learning-based reconstruction and uses only off-the-shelf components without relay optics or cascaded modules. We select the birefringent filter configuration using a condition-number-based criterion and validate the system on both synthetic and real data.

Keywords: Snapshot Hyperspectral Imaging · Angular-to-Spectral Diversity Conversion · Kaleidoscope · Birefringent Filter

1 Introduction

Hyperspectral imaging (HSI) captures 3D spectral data cubes, providing much denser spectral information than the human eye or an RGB camera. HSI has been widely used in computer vision tasks such as object recognition [22], color consistency [1], and anomaly detection [25]. Conventionally, HSI is acquired by scanning, sequentially capturing 2D slices and incurring long capture times and complex systems.

In contrast, snapshot HSI [7, 11] reconstructs a spectral data cube from a single exposure, avoiding sequential scanning and simplifying acquisition. However, snapshot HSI systems that simultaneously achieve high reconstruction performance, low cost, and compact optics remain limited, which hinders broad adoption. Existing methods often satisfy only a subset of these three requirements.

[†] Corresponding author.




Fig. 1: A kaleidoscope produces angularly diverse replicated views, and an attached BRF converts them into spectral-response diversity for snapshot HSI.

Approaches based on diffractive optical elements (DOEs) [13, 24] or custom filter arrays [5, 30] require custom fabrication or permanent hardware modification. They can be compact and achieve high reconstruction performance, but their manufacturing cost is high. By contrast, methods using an RGB camera, with or without a single additional optical element (e.g., a prism) [2, 4, 9, 20, 32], suffer from weak or spatially localized spectral modulation, which makes the inverse problem ill-conditioned. They can be low-cost and compact, but challenges remain in reconstruction accuracy or computational cost. Spatial-replication HSI [10, 14, 18] often enables simple and fast reconstruction, but image replication and spectral modulation are commonly implemented in separate stages, which increases system size due to relay optics or cascaded components.

To address this gap, we propose a spatial-replication snapshot HSI system that converts angular diversity into spectral-response diversity, as illustrated in Fig. 1. A tapered kaleidoscope generates multiple replicated views with distinct incidence directions, and a directly attached birefringent filter (BRF) translates these directional differences into channel-dependent spectral transmittances. Rather than treating the BRF’s incidence-direction dependence as an error source, as is commonly done in conventional BRF-based imaging systems [10, 15, 16, 23], we deliberately leverage it as a source of spectral-response diversity. As a result, the replicated views provide complementary measurements that better condition the inverse problem for stable single-shot spectral reconstruction. The proposed design also preserves a simple pixel-wise linear observation model, enabling fast and stable non-learning-based reconstruction. It remains low-cost by relying on off-the-shelf optical components, and compact by integrating image replication and spectral modulation into a single attached unit that eliminates relay optics and cascaded modules.

The main contributions of this paper are as follows:

- We propose angular-to-spectral diversity conversion for snapshot HSI by using kaleidoscope multi-view sampling and the incidence-direction dependence of a BRF, with a condition-number-based BRF selection criterion.
- We present a low-cost ($\sim \$1k$) and compact 21 cm off-the-shelf kaleidoscope-BRF add-on that directly integrates spatial replication and BRF spectral modulation, eliminating relay stages.

2 Related Work

Snapshot HSI for Low-Cost and Compact Systems. Low-cost and compact snapshot HSI has been explored in several forms. RGB-only approaches [2, 20, 32] are attractive for their simplicity, but their spectral responses are often limited and highly correlated, which can make the inverse problem ill-conditioned. Adding a single coding element can improve measurement diversity. However, the resulting spectral cues often remain weak or spatially localized. Baek *et al.* [4] combine an RGB camera with a prism, where spectral cues are sparse and reconstruction is solved iteratively with a reported runtime of about 45 minutes. Zhao *et al.* [33] use a printer-generated random mask and reconstruct with a CNN trained on large paired data, with training reportedly taking about 24 hours on a single GPU. Oh *et al.* [21] realize HSI using three cameras with different spectral sensitivity profiles, but this requires synchronized multi-camera capture and increases system size.

Low-Cost Spatial-Replication Snapshot HSI. Spatial-replication snapshot HSI is closest to our approach. We focus on low-cost designs, rather than higher-end alternatives [10, 19]. Manakov *et al.* [18] use a kaleidoscope purely as an image multiplier, then combine a filter array with relay optics so that each sub-image receives a distinct spectral modulation. A key drawback is that the relay optics increase the system size. The overall system length is reported to be approximately 100 cm [18]. Takatani *et al.* [26] attach color filters directly onto mirror surfaces to simultaneously realize image replication and spectral multiplexing without relay optics. However, the filter-on-mirror configuration can introduce diffuse reflections at the added interfaces, which may lead to blur in the replicated sub-images. Compared with these systems, our method also uses spatial replication, but realizes spectral diversity through an attached BRF and incidence-direction differences, without relay optics or filter-on-mirror designs.

BRF-Based Spectral Imaging. BRFs have been used for spectral imaging by varying the optical path difference (OPD) over time, for example by using liquid crystal controlled by an applied voltage [3, 8, 34] or by mechanically rotating polarizers or birefringent elements [23]. Such approaches obtain measurement diversity through time multiplexing and are therefore scanning HSI methods. Snapshot HSI with birefringent optics includes cascaded polarization-splitting designs [10], in which increasing the number of channels generally requires a deeper cascade, as well as spatially varying OPD in natural materials [9], where spectral cues can remain localized. Another line of work uses birefringent interferometry [15], which typically relies on specialized optics. In contrast, our

method obtains channel diversity from incidence-direction differences using a single attached BRF together with a kaleidoscope.

Oblique-Incidence Effects in BRF. Because BRF transmittance varies with incidence direction [16, 27, 31], oblique incidence in imaging is often treated as a source of spectral non-uniformity. In HSI, pixel-dependent ray angles arise naturally from the field of view and finite aperture cone, so many analyses assume near-normal incidence or adopt small-angle approximations. Other studies explicitly model this effect or mitigate it by limiting the incidence cone or compensating wavelength shifts in post-processing [16, 17, 31]. By contrast, in laser systems, BRFs are often operated under oblique incidence, where the incidence angle and optic axis parameters are deliberately chosen to control spectral selection and improve laser performance [6, 28, 29]. Our method follows this latter perspective. Rather than suppressing angle dependence, we exploit it as a source of diverse channel-wise spectral modulations for snapshot HSI.

A broader qualitative comparison with representative snapshot HSI systems is provided in Supp. Mat. Sec. 2.

3 A Pixel-Wise Linear Measurement Model for HSI

We use a standard pixel-wise linear model for snapshot HSI, assuming that all measurement channels are registered to the same scene point. Let $m \in \{1, \dots, M\}$ index the measurements, and let $T_m(\lambda)$ denote the effective spectral response of channel m . The m -th measurement is modeled as

$$I_m = \int_{\Lambda} l(\lambda) T_m(\lambda) s(\lambda) d\lambda, \quad (1)$$

where $l(\lambda)$ is the illumination spectrum and $s(\lambda)$ is the spectral reflectance.

Discretizing Equation (1) over sampled wavelengths $\{\lambda_k\}_{k=1}^{N_\lambda}$ gives

$$\mathbf{i} = \mathbf{H}\mathbf{s}, \quad (2)$$

where $\mathbf{i} \in \mathbb{R}^M$ is the measurement vector, $\mathbf{s} \in \mathbb{R}^{N_\lambda}$ is the spectral reflectance, and $\mathbf{H} \in \mathbb{R}^{M \times N_\lambda}$ is the sensing matrix with

$$H_{m,k} = l(\lambda_k) T_m(\lambda_k) \Delta\lambda_k. \quad (3)$$

Reconstruction is governed by $\mathbf{H}$, and its stability depends on how distinct the channel-wise spectral responses $\{T_m(\lambda)\}_{m=1}^M$ are.

4 Proposed System

Building on the discussion in Sec. 3, this section introduces the proposed system and explains how it converts angular diversity into spectral-response diversity.




Fig. 2: Proposed optics: (a) System overview; (b) angular-to-spectral diversity conversion by the off-surface optic axis BRF; (c) angular geometry in the BRF.

4.1 System Overview

To clarify the core idea behind our method, we begin with an overview of the optical system shown in Figures 2a and 2b. The system consists of an objective lens, a diffuser, a BRF, a tapered kaleidoscope, and a camera.

The objective lens forms a 2D image of the scene on the diffuser plane. The diffuser suppresses parallax induced by scene depth, ensuring that the replicated images share the same spatial structure across viewpoints [14, 18]. However, the diffuser's Bidirectional Transmittance Distribution Function (BTDF) introduces differences in the spectral intensity distribution. This issue is addressed through calibration in our method.

A BRF is placed immediately behind the diffuser. The BRF is mounted at the entrance of the tapered kaleidoscope. The tapered kaleidoscope is a hollow mirrored tube that replicates the input image into multiple sub-images through internal reflections [12]. These replicated sub-images are recorded by a single camera.

4.2 Angular-to-Spectral Diversity Conversion

The central mechanism of the proposed system is angular-to-spectral diversity conversion. The system first generates diversity in incidence directions across view channels, and then converts that angular diversity into spectral-modulation diversity through the incidence-direction-dependent spectral modulation of the BRF.

The tapered kaleidoscope produces multiple replicated sub-images, which we call view channels. Because these channels correspond to different reflection paths, they also correspond to different incidence directions at the BRF. Let (α_v, σ_v) denote the incidence direction associated with view channel $v \in \{1, \dots, V\}$, where α_v is the incidence angle and σ_v is the azimuth (Figure 2c). The taper enlarges the spread of these sampled directions, increasing the angular diversity across channels.

This angular diversity is then converted into channel-wise spectral-response diversity by the BRF. A BRF is implemented as a birefringent plate sandwiched between two linear polarizers. After passing through the first polarizer, the incident light becomes linearly polarized and splits inside the birefringent plate into ordinary and extraordinary components. These components accumulate an OPD, denoted by Δ_{o-e} , and the second polarizer projects both components onto a common polarization axis, resulting in a transmitted intensity that varies with Δ_{o-e}/λ . Because Δ_{o-e} depends on the incidence direction, the BRF spectral transmittance also changes with (α, σ) . Under the crossed-Nicols configuration shown in Figure 2c, the BRF spectral transmittance is written as

$$B(\lambda; \alpha, \sigma, \theta) \propto \sin^2 \left(\frac{\pi \Delta_{o-e}(\alpha, \sigma, \theta)}{\lambda} \right), \quad (4)$$

where θ denotes the optic axis diving angle (the angle between the optic axis and the surface plane). Therefore, the different incidence directions $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ generated by the kaleidoscope are mapped to different BRF spectral transmittances,

$$B_v(\lambda) = B(\lambda; \alpha_v, \sigma_v, \theta). \quad (5)$$

The camera further observes each view through one of the RGB sensor sensitivities $f^c(\lambda)$, where $c \in \{R, G, B\}$. Hence, the effective spectral response associated with view channel v and color channel c becomes

$$T_{v,c}(\lambda) = f^c(\lambda) B_v(\lambda). \quad (6)$$

As a result, the angular diversity produced by the kaleidoscope is converted into channel-wise spectral-response diversity $\{\{T_{v,c}(\lambda)\}_{v=1}^V\}_{c \in \{R, G, B\}}$, which provides the measurement diversity required for stable spectral reconstruction.

Thanks to the diffuser, the view channels retain nearly the same spatial structure and differ mainly in spectral modulation. This enables snapshot spectral multiplexing with a simple channel-wise image formation model. The kaleidoscope produces angular diversity, and the BRF converts it into spectral-response diversity across channels, yielding complementary measurements that improve the conditioning of the inverse problem. Because the BRF is mounted directly at the entrance of the tapered kaleidoscope, this angular-to-spectral conversion is realized within a compact optical architecture, without relay optics or cascaded components.




Fig. 3: Simulated OPD magnitude as a function of incidence direction (α, σ) . Top: 3D surface with $(x, y) = (\alpha \cos \sigma, \alpha \sin \sigma)$. Bottom: polar contour with radial coordinate α and azimuth σ . Parameters: $n = 1$, $n_o = 1.54264$, $n_e = 1.55170$, $L = 5 \times 10^{-4}$ m.

4.3 Pixel-Wise Spectral Reconstruction

Using the effective spectral response $T_{v,c}(\lambda)$ defined in Equation (6), the sensing matrix of the proposed system is given by

$$H_{(v,c),k} = l(\lambda_k) f^c(\lambda_k) B_v(\lambda_k) \Delta\lambda_k. \quad (7)$$

After inter-channel registration, we reconstruct the spectral reflectance at each pixel by solving a non-negative least-squares problem with second-order spectral smoothness:

$$\hat{\mathbf{s}}(\mathbf{x}) = \arg \min_{\mathbf{s} \geq 0} \|\mathbf{H}\mathbf{s} - \mathbf{i}(\mathbf{x})\|_2^2 + \beta \|\mathbf{D}\mathbf{s}\|_2^2, \quad (8)$$

where $\mathbf{D}$ is the second-order difference matrix along wavelength, and β is the smoothness weight.

Because the proposed system provides sufficient measurement diversity through diverse channel-wise spectral responses while preserving a simple spatial-replication image formation model, spectral reconstruction reduces to a lightweight regularized linear inverse problem at each pixel, without learning.

5 BRF Design and Parameter Selection

This section presents a BRF design procedure for the proposed system, using a conditioning-based criterion to select a suitable off-the-shelf configuration.

5.1 OPD-Based Model Under Oblique Incidence

In the proposed system, the tapered kaleidoscope determines a discrete set of incidence directions $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ at the BRF. For a birefringent plate under oblique incidence with a diving optic axis, the OPD can be written as follows [27]:

$$\Delta_{o-e} = L \left[\sqrt{n_o^2 - (n \sin \alpha)^2} + \frac{n(n_o^2 - n_e^2) \sin \theta \cos \theta \cos \sigma \sin \alpha}{D} - \frac{n_o}{D} \sqrt{n_e^2 D - (n_e^2 - (n_e^2 - n_o^2) \cos^2 \theta \sin^2 \sigma)(n \sin \alpha)^2} \right], \quad (9)$$

where $D = n_e^2 \sin^2 \theta + n_o^2 \cos^2 \theta$.

Here, n is the refractive index of the surrounding isotropic medium, n_o and n_e are the ordinary and extraordinary refractive indices of the birefringent plate, and L is the plate thickness.

Combined with Equation (4), Equation (9) determines how the BRF spectral transmittance varies across the sampled incidence directions. Accordingly, the BRF design problem is to choose θ such that the resulting set of channel-wise transmittances $\{B_v(\lambda)\}_{v=1}^V$ provides strong spectral-response diversity.

5.2 Design Implications of OPD Symmetry

A large variation of Δ_{o-e} over (α, σ) is helpful, but it is not sufficient by itself. What matters is whether the sampled directions $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ map to distinct OPD values. If multiple view channels yield similar OPDs, their BRF transmittances become similar, the corresponding measurements are redundant, and the effective spectral responses in $\mathbf{H}$ become more correlated.

This issue is particularly important for spatial-replication snapshot HSI, where the sampled directions are constrained by the kaleidoscope geometry and often contain structured azimuthal relationships. As an extreme case, when $\theta = 90^\circ$, Equation (9) becomes independent of σ , so the BRF provides no azimuthal diversity.

In the conventional on-surface optic axis BRF configuration ($\theta = 0^\circ$), the OPD is invariant under a 180° azimuthal rotation,

$$\Delta_{o-e}(\alpha, \sigma; 0) = \Delta_{o-e}(\alpha, \sigma + \pi; 0), \quad (10)$$

which causes channel pairs under symmetric angular sampling to collapse to the same OPD. This reduces measurement diversity by making the channel-wise spectral modulations more similar and can make the inverse problem ill-conditioned.

By contrast, for an off-surface optic axis BRF ($0 < \theta < 90^\circ$), the term proportional to $\cos \sigma \sin \alpha$ in Equation (9) breaks the $\sigma \rightarrow \sigma + \pi$ invariance. As shown in Figure 3, this removes the half-turn repetition observed in the on-surface optic axis case and reduces redundancy among the sampled view channels. A mirror symmetry still remains,

$$\Delta_{o-e}(\alpha, \sigma; \theta) = \Delta_{o-e}(\alpha, -\sigma; \theta), \quad (11)$$

but this pairing is typically less harmful than the $\sigma \rightarrow \sigma + \pi$ equivalence under rotationally structured sampling.

Moreover, Figure 3 shows that off-surface optic axis settings also increase the spread of $|\Delta_{o-e}|$ over (α, σ) compared with the on-surface optic axis case. This broader OPD distribution strengthens the differences among the channel-wise BRF transmittances in Equation (4). Together, these observations motivate the use of an off-surface optic axis BRF in the proposed system.

5.3 Design Criterion: Conditioning of the Sensing Matrix

The remaining question is which off-surface optic axis configuration best improves spectral reconstruction. Because the effective spectral responses are encoded in the sensing matrix $\mathbf{H}$ in Equation (7), we evaluate BRF designs through the conditioning of $\mathbf{H}$.

Specifically, we use the condition number

$$\kappa(\mathbf{H}) := \frac{\sigma_{\max}(\mathbf{H})}{\sigma_{\min}(\mathbf{H})}, \quad (12)$$

where $\sigma_{\max}$ and $\sigma_{\min}$ are the largest and smallest singular values of $\mathbf{H}$, respectively. When $\sigma_{\min}$ is small, $\mathbf{H}$ becomes nearly rank-deficient, the inverse problem becomes unstable, and stronger regularization is required. Therefore, a smaller $\kappa(\mathbf{H})$ indicates that the channel-wise spectral modulations are more complementary and that reconstruction is more stable.

In our setting, the incidence-direction samples $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ are determined by the tapered-kaleidoscope geometry, while $l(\lambda)$ and the camera sensitivities $f^c(\lambda)$ are obtained by calibration. Hence, once the sampled directions are fixed, we can precompute $\mathbf{H}$ for each candidate BRF configuration and use $\kappa(\mathbf{H})$ as a practical design metric. This gives a direct parameter-selection procedure: determine the angular samples produced by the kaleidoscope, evaluate candidate BRF settings through $\kappa(\mathbf{H})$, and choose the configuration that yields the best conditioning. We next apply this procedure to the physical parameter choices used in our prototype.

An experimental validation of the OPD distribution using measured transmittance profiles is provided in Supp. Mat. Sec. 4.

5.4 Applying the Criterion to Physical Parameter Selection

We first fix the tapered-kaleidoscope geometry, which defines the discrete incidence-direction samples at the BRF, and then select the BRF configuration that best conditions the sensing matrix under those samples.

Tapered kaleidoscope. We use a tapered kaleidoscope to obtain a diverse set of incidence directions at the BRF. Following [12], we adopt a 3-sided design and set the taper angle to 9° , which reduces view-channel fragmentation on the sensor while maintaining sharp and stable multi-reflection views.

Because tapering makes the replicated virtual images behave as samples on a virtual viewing sphere, different reflection orders can have different apparent




Fig. 4: Optic axis selection using conditioning.

distances. To keep all view channels simultaneously in focus without requiring an excessively small aperture, we use the $V = 10$ view channels produced by reflection orders up to two. Given this kaleidoscope geometry, we compute the sampled incidence directions $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ by ray tracing, and we also scan the in-plane mounting rotation b , which shifts the azimuths as $\sigma_v \leftarrow \sigma_v + b$. These discrete angular samples are then used in the conditioning-based BRF selection below.

BRF. Given the sampled incidence directions $\{(\alpha_v, \sigma_v)\}_{v=1}^V$ obtained above, we select the BRF by minimizing the condition number introduced in Section 5.3. For each candidate optic axis diving angle θ , we scan the in-plane mounting rotation b and evaluate the resulting $\kappa(\mathbf{H})$. Figure 4a summarizes the results.

We use a synthetic quartz plate as the birefringent plate in the BRF. Synthetic quartz is widely used in frequency-control resonators and is commercially available in standard cuts with known optic axis angles. This makes it a practical low-cost choice that is consistent with our off-the-shelf design goal. We therefore sweep these cut-defined candidates and select the one that best conditions $\mathbf{H}$. Among these candidates, an ST-cut synthetic quartz plate ($\theta = 47.25^\circ$) yields the best conditioning. We thus use an ST-cut synthetic quartz plate (MTI Corporation, $10\text{ mm} \times 10\text{ mm} \times 0.5\text{ mm}$) and assemble a BRF in a crossed-Nicols configuration using two linear film polarizers.

For the in-plane mounting rotation b , the condition number is exceptionally large at $b = 0^\circ$ (1.66×10^{16}), indicating that $\mathbf{H}$ is effectively rank-deficient. This is consistent with the residual symmetry of the off-surface optic axis OPD map (Fig. 4b). At this rotation, multiple sampled incidence directions fall onto nearly identical OPD values, which reduces spectral-modulation diversity. We therefore avoid $b = 0^\circ$ and assemble the system at $b = 20^\circ$.

6 Implementation Details

6.1 Other Hardware Setup

Diffuser. We use a wide-angle holographic diffuser (Edmund Optics, 80° Difusing Angle, 2 inch square) to support the wide range of incidence angles in our multi-view observation. Its transmission drops to about 50% at an incidence angle of 40°. Given our maximum incidence angle of 34.86°, this angular roll-off can reduce the optical throughput of high-reflection view channels.

Other physical dimensions. The prototype is a 21 cm add-on unit (approx. 40 cm including the camera) built with 35 mm objective lens (Edmund Optics, C-mount) and a camera (Sony α7R IV) with a zoom lens (TAMRON, at the telephoto end: 40 mm). The camera is placed approximately 30 cm from the BRF, constrained by the lens minimum focus distance.

Add-on component parts cost. The approximate costs of the add-on components are: polarizers ≈\$8, mirrors ≈\$20 total (custom-cut front-surface mirrors readily ordered from a consumer-facing retailer), synthetic quartz \$33.35, diffuser \$273, and objective lens \$635. Overall, the deployed add-on hardware remains within \$1k, excluding the camera body, imaging lens, and calibration equipment such as a monochromator.

6.2 Preprocessing

To construct the sensing matrix $\mathbf{H}$ in Equation (7), we pre-estimate (i) the view-channel-wise BRF transmittance $B_v(\lambda)$, (ii) the product of the illumination spectrum and camera RGB spectral sensitivity, $\phi^c(\lambda) = l(\lambda)f^c(\lambda)$, and wavelength-independent throughput factors g_v .

(i) Measuring $B_v(\lambda)$. We measure $B_v(\lambda)$ using monochromator-based narrow-band illumination and estimate the channel-wise transmittance profiles from the captured images. This calibration absorbs the diffuser BTDF into the estimated $B_v(\lambda)$. To account for slight within-channel variation, we use pixel-wise calibrated transmittance profiles when constructing the per-pixel sensing matrix.

(ii) Estimating $\phi^c(\lambda)$ and throughput factors g_v . Assuming a spectrally flat reflectance for a white reference target, the observation of view channel v and color channel c under the white target can be written as

$$I_{v,w}^c \approx g_v \int_{\Lambda} \phi^c(\lambda) B_v(\lambda) d\lambda, \quad (13)$$

where g_v accounts for channel-wise throughput differences. We discretize Equation (13) and estimate $\phi^c(\lambda)$ and g_v jointly via regularized least squares. We then apply channel-wise gain correction and construct $\mathbf{H}$ from the calibrated quantities according to Equation (7).

Inter-channel registration. We align view channels using a calibration target and estimate inter-channel homographies, including mirror-induced parity flips. The estimated registration parameters, $\mathbf{H}$, and g_v are fixed in subsequent reconstruction.

Full calibration details are provided in Supp. Mat. Sec. 1.




Fig. 5: Synthetic ColorChecker comparison with low-cost snapshot baselines.

7 Experiments

7.1 Experiments on Synthetic Data

Experimental Setup. We evaluate our system against two low-cost snapshot baselines, *3Camera* [21] and *Faced Reflectors* [26], focusing on the measurement diversity induced by our channel-wise spectral-response diversity. To isolate the effect of the optics, we reconstruct spectra for all methods using the same non-negative, smoothness-regularized least-squares formulation in Equation (8). We generate synthetic observations from the 24 ColorChecker reflectance spectra using each method’s image formulation model, and add i.i.d. Gaussian noise after exposure matching such that the maximum DN is approximately 255. The measured transmittance profiles used in our method are shown in Figure 5b. For each noise level, we sweep the smoothness weight β and report the best PSNR over β , as shown in Figure 5a.

See Supp. Mat. Sec. 3 for details.

Result Analysis. Our method achieves higher PSNR than the baselines in the low-noise regime, suggesting that the proposed system provides richer spectral-response diversity across channels and hence greater measurement diversity for the inverse problem. As noise increases, the PSNR gap narrows and can slightly reverse in the high-noise regime. This behavior may be explained by non-uniform optical throughput across view channels. Because our channels differ not only in spectral transmittance but also in overall optical throughput, the resulting SNR is imbalanced; when noise dominates, low-throughput channels contribute less as effective measurements, diminishing the benefit of additional channels. This throughput imbalance is amplified by angle-dependent diffuser emission and by the polarizer loss in the BRF.

7.2 Experiments on Real Data

ColorChecker Validation. We calibrate the illumination spectrum and camera RGB sensitivities using a gray balance target, and reconstruct spectral reflectance of an X-Rite ColorChecker Passport Photo 2 from the 10 view channels




Fig. 6: ColorChecker reflectance reconstruction: (a) capture crop; (b) RGB rendering; red dots mark the patch sampling locations; (c) PSNR vs. smoothness weight β ; (d) reconstructed and reference spectra.

in Figure 6a. We illuminate the scene with an incandescent lamp. We reconstruct 25 bands over 430 nm to 670 nm at 10 nm intervals. We used conservative settings (ISO 125, F/8, 8 s) for robust focus and SNR across view channels. Since the recovered reflectance is defined up to a global scale, we fit a single scalar using all 24 patches by least squares against the reference spectra and apply it to all reconstructions. We determine a single smoothness weight β from the ColorChecker experiment (Figure 6c) and keep it fixed for the real-scene and USAF-target reconstructions. At this β , Figures 6b and 6d show the RGB rendering and reconstructed spectra. The RGB rendering is obtained by projecting the reconstructed reflectance cube onto the calibrated camera RGB sensitivities $f^c(\lambda)$, $c \in \{R, G, B\}$. Each rendered channel is normalized by the 95th percentile of valid pixels, and no display gamma correction is applied. The PSNR computed over all 24 patches and all wavelengths is 27.4 dB, confirming that our system enables reasonably accurate spectral reconstruction.

Runtime. We measure per-frame reconstruction time for 25-band reflectance on 226×155 pixels, averaged over five RAW frames, on a workstation with an NVIDIA Quadro RTX 8000 GPU. The total runtime is 0.748 s/frame, of which solving Equation (8) takes 0.044 s/frame ($\approx 5.9\%$), while most time is spent on RAW loading and forming the observation vector $\mathbf{i}$. A one-time preprocessing




Fig. 7: Real-scene 25-band reconstruction. Colored outlines indicate registered view-channel regions; the white dashed box is used for the valid-area ratio.

step for illumination and camera-sensitivity calibration takes ~ 26 s and is reused as long as the configuration is unchanged.

Real-Scene Validation. To evaluate the prototype beyond flat ColorChecker patches, we captured a real scene containing a leaf, a flower, and printed materials under daylight-level artificial solar illumination (SERIC SOLAX XC-100). The scene was captured using the same prototype with a ~~0.25 s exposure, ISO 200, and F/8~~. We reconstructed 25 bands over 430 nm to 670 nm using the same calibration, inter-channel registration, reconstruction pipeline, and the β selected in the ColorChecker validation. As shown in Figure 7, the reconstructed spectral bands and point spectra exhibit material-dependent spectral variations across the scene. Since the commercial scanning-HSI reference was captured from a different viewpoint, it is not pixel-aligned with our reconstruction and is used only to compare spectral trends at the marked points. Residual errors remain near object boundaries and printed-text edges, indicating the practical limits of registration accuracy and diffuser-based parallax suppression.

Spatial Resolution. To quantify the spatial-resolution trade-off of the 10-view configuration, we evaluated a USAF target. Full details are provided in Supp. Mat. Sec. 5. Most reconstructed HSI bands show resolvable structures at 0.89 lp/mm to 1.00 lp/mm, compared with 1.12 lp/mm for an RGB capture without the kaleidoscope/diffuser under the same objective lens, camera, and image-plane scale. Thus, the limiting resolution is within one to two USAF group/element steps of the RGB capture without the kaleidoscope/diffuser. The main spatial cost of the current prototype is instead the reduced valid reconstruction area caused by 10-view replication and the non-optimized camera-kaleidoscope field-of-view matching. The valid region contains approximately 3.24×10^5 pixels, or 4.5 % of the dashed box enclosing the view-channel regions in Figure 7.

7.3 Ablation Study

We ablate the design factors that determine measurement diversity under the same synthetic protocol as in Section 7.1. The goal is to test whether reconstruction accuracy follows the conditioning-based design analysis in Section 5.




Fig. 8: Synthetic ColorChecker ablation study.

Effect of the Number of View Channels. We vary the number of used view channels from $V = 10$ down to 1. PSNR decreases monotonically as V decreases (Figure 8a), showing that the view channels provide non-redundant spectral constraints for reconstruction.

Effect of BRF Angle Configuration. We vary the optic axis angle θ and the mounting rotation b , and evaluate PSNR with the same reconstructor. As shown in Figure 8b, on-surface optic axis or poorly aligned settings are consistently worse than the selected off-surface optic axis setting, indicating that the BRF angle configuration directly affects usable spectral diversity.

Relationship Between Conditioning and PSNR. Figure 8c shows that PSNR follows the trend of $\kappa(\mathbf{H})$ across (θ, b) , supporting condition-number minimization as a practical design rule.

8 Conclusion

We presented a compact, low-cost spatial-replication snapshot HSI system that uses a directly attached tapered kaleidoscope–BRF unit to convert kaleidoscope-induced angular diversity into BRF-based spectral-response diversity. The proposed design preserves a pixel-wise linear measurement model, enables fast non-learning-based spectral reconstruction, and selects the BRF configuration using a condition-number-based criterion. Experiments on synthetic data, ColorChecker validation, and real-scene reconstruction with an off-the-shelf prototype validate stable and reasonably accurate 25-band reconstruction in a compact add-on form factor.

Limitations. The prototype exhibits channel-dependent throughput due to diffruser angular roll-off, which reduces the effective contribution of some view channels in noise-dominated regimes and leads to long exposure under conservative ColorChecker settings. Its effective spatial resolution is limited by dividing the sensor area among replicated views. The current implementation also requires one-time monochromator-based calibration outside the add-on component cost; reducing this burden through model-assisted calibration is future work.

Acknowledgements This work was supported by JST SPRING JPMJSP2140, JSPS KAKENHI JP23H00499, and JST CREST JPMJCR23N3.

References

1. Abrardo, A., Alparone, L., Cappellini, I., Prosperi, A.: Color constancy from multispectral images. In: Proceedings 1999 International Conference on Image Processing (Cat. 99CH36348). vol. 3, pp. 570–574 vol.3 (1999). <https://doi.org/10.1109/ICIP.1999.817179>
2. Arad, B., Ben-Shahar, O.: Sparse recovery of hyperspectral signal from natural rgb images. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) Computer Vision – ECCV 2016. pp. 19–34. Springer International Publishing, Cham (2016)
3. August, I., Oiknine, Y., AbuLeil, M., Abdulhalim, I., Stern, A.: Miniature compressive ultra-spectral imaging system utilizing a single liquid crystal phase retarder. Scientific reports **6**(1), 1–9 (2016)
4. Baek, S.H., Kim, I., Gutierrez, D., Kim, M.H.: Compact single-shot hyperspectral imaging using a prism. ACM Trans. Graph. **36**(6) (Nov 2017). <https://doi.org/10.1145/3130800.3130896>
5. Bian, L., Wang, Z., Zhang, Y., Li, L., Zhang, Y., Yang, C., Fang, W., Zhao, J., Zhu, C., Meng, Q., Peng, X., Zhang, J.: A broadband hyperspectral image sensor with high spatio-temporal resolution. Nature **635**(8037), 73–81 (Nov 2024). <https://doi.org/10.1038/s41586-024-08109-1>
6. Demirbas, U.: Off-surface optic axis birefringent filters for smooth tuning of broadband lasers. Appl. Opt. **56**(28), 7815–7825 (Oct 2017). <https://doi.org/10.1364/AO.56.007815>
7. Ding, K., Wang, M., Chen, M., Wang, X., Ni, K., Zhou, Q., Bai, B.: Snapshot spectral imaging: from spatial-spectral mapping to metasurface-based imaging. Nanophotonics **13**(8), 1303–1330 (2024). <https://doi.org/doi:10.1515/nanoph-2023-0867>
8. Fan, A., Xu, T., Li, J., Teng, G., Wang, X., Zhang, Y., Xu, C.: Compressive full-stokes polarization and flexible hyperspectral imaging with efficient reconstruction. Optics and Lasers in Engineering **160**, 107256 (2023). <https://doi.org/10.1016/j.optlaseng.2022.107256>
9. Fujimoto, Y., Kitano, K., Fujiwara, K., Fujimura, Y., Funatomi, T., Mukaigawa, Y.: Snapshot hyperspectral imaging using petrographic thin section. In: International Conference on Machine Vision Applications (MVA) 2025 (Jul 2025). <https://doi.org/10.34385/proc.93.P2-4>, session P2-4, Publication date: 2025-07-26
10. Gorman, A., Fletcher-Holmes, D.W., Harvey, A.R.: Generalization of the lyot filter and its application to snapshot spectral imaging. Opt. Express **18**(6), 5602–5608 (Mar 2010). <https://doi.org/10.1364/OE.18.005602>
11. Hagen, N.A., Kudenov, M.W.: Review of snapshot spectral imaging technologies. Optical Engineering **52**(9), 090901 (2013). <https://doi.org/10.1117/1.OE.52.9.090901>
12. Han, J.Y., Perlin, K.: Measuring bidirectional texture reflectance with a kaleidoscope. ACM Trans. Graph. **22**(3), 741–748 (Jul 2003). <https://doi.org/10.1145/882262.882341>
13. Jeon, D.S., Baek, S.H., Yi, S., Fu, Q., Dun, X., Heidrich, W., Kim, M.H.: Compact snapshot hyperspectral imaging with diffracted rotation. ACM Transactions on Graphics (Proc. SIGGRAPH 2019) **38**(4), 117:1–13 (2019). <https://doi.org/10.1145/3306346.3322946>
14. Jeon, D.S., Choi, I., Kim, M.H.: Multisampling compressive video spectroscopy. In: Proceedings of the 37th Annual Conference of the European Association for

- Computer Graphics. p. 467–477. EG '16, Eurographics Association, Goslar, DEU (2016)
15. Kudenov, M.W., Dereniak, E.L.: Compact real-time birefringent imaging spectrometer. *Opt. Express* **20**(16), 17973–17986 (Jul 2012). <https://doi.org/10.1364/OE.20.017973>
 16. Li, X., Dong, K., Cao, Z., An, Y., Zhang, L., Wang, H., Zheng, X.: Investigation of the field-of-view properties in lyot liquid crystal tunable filter. *Opt. Express* **33**(26), 54112–54127 (Dec 2025). <https://doi.org/10.1364/OE.583121>
 17. Mahajan, M., Murray, B., Gu, D.F., Zhang, W., Mansell, J.: Voltage tunable multi-spectral imaging lens. In: Dolne, J.J., Spencer, M.F. (eds.) *Unconventional Imaging and Adaptive Optics 2021*. vol. 11836, p. 118360B. International Society for Optics and Photonics, SPIE (2021). <https://doi.org/10.1117/12.2594854>
 18. Manakov, A., Restrepo, J.F., Klehm, O., Hegedüs, R., Eisemann, E., Seidel, H.P., Ihrke, I.: A reconfigurable camera add-on for high dynamic range, multispectral, polarization, and light-field imaging. *ACM Trans. Graph.* **32**(4) (Jul 2013). <https://doi.org/10.1145/2461912.2461937>
 19. Mu, T., Han, F., Bao, D., Zhang, C., Liang, R.: Compact snapshot optically replicating and remapping imaging spectrometer (orris) using a focal plane continuous variable filter. *Opt. Lett.* **44**(5), 1281–1284 (Mar 2019). <https://doi.org/10.1364/OL.44.001281>
 20. Nguyen, R.M.H., Prasad, D.K., Brown, M.S.: Training-based spectral reconstruction from a single rgb image. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *Computer Vision – ECCV 2014*. pp. 186–201. Springer International Publishing, Cham (2014)
 21. Oh, S.W., Brown, M.S., Pollefeys, M., Kim, S.J.: Do it yourself hyperspectral imaging with everyday digital cameras. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2461–2469 (2016). <https://doi.org/10.1109/CVPR.2016.270>
 22. Okajima, R., Tominaga, S.: Object Recognition by Multi-Spectral Imaging with a Liquid Crystal Filter . In: *Pattern Recognition, International Conference on*. vol. 1, p. 1708. IEEE Computer Society, Los Alamitos, CA, USA (Sep 2000). <https://doi.org/10.1109/ICPR.2000.905485>
 23. Salesin, K., Seyb, D., Friday, S., Jarosz, W.: Diy hyperspectral imaging via polarization-induced spectral filters. In: *2022 IEEE International Conference on Computational Photography (ICCP)*. pp. 1–12. IEEE (2022)
 24. Shi, Z., Chugunov, I., Bijelic, M., Côté, G., Yeom, J., Fu, Q., Amata, H., Heidrich, W., Heide, F.: Split-aperture 2-in-1 computational cameras. *ACM Trans. Graph.* **43**(4) (Jul 2024). <https://doi.org/10.1145/3658225>
 25. Stein, D., Beaven, S., Hoff, L., Winter, E., Schaum, A., Stocker, A.: Anomaly detection from hyperspectral imagery. *IEEE Signal Processing Magazine* **19**(1), 58–69 (2002). <https://doi.org/10.1109/79.974730>
 26. Takatani, T., Aoto, T., Mukaigawa, Y.: One-shot hyperspectral imaging using faced reflectors. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4039–4047 (2017)
 27. Veiras, F.E., Perez, L.I., Garea, M.T.: Phase shift formulas in uniaxial media: an application to waveplates. *Appl. Opt.* **49**(15), 2769–2777 (May 2010). <https://doi.org/10.1364/AO.49.002769>
 28. Wei, J., Cao, X., Jin, P., Su, J., Lu, H., Peng, K.: Diving angle optimization of brf in a single-frequency continuous-wave wideband tunable titanium:sapphire laser. *Opt. Express* **29**(5), 6714–6725 (Mar 2021). <https://doi.org/10.1364/OE.419580>

29. Wei, J., Su, J., Lu, H., Peng, K.: A review of progress about birefringent filter design and application in ti:sapphire laser. *Photonics* **10**(11) (2023). <https://doi.org/10.3390/photonics10111217>
30. Yako, M., Yamaoka, Y., Kiyohara, T., Hosokawa, C., Noda, A., Tack, K., Spooren, N., Hirasawa, T., Ishikawa, A.: Video-rate hyperspectral camera based on a cmos-compatible random array of fabry-pérot filters. *Nature Photonics* **17**(3), 218–223 (Mar 2023). <https://doi.org/10.1038/s41566-022-01141-5>
31. Zhang, J., Li, J.: Spectral shift correction and adaptive band selection for multi-spectral imaging. *Optics and Lasers in Engineering* **144**, 106632 (2021). <https://doi.org/10.1016/j.optlaseng.2021.106632>
32. Zhang, J., Su, R., Fu, Q., Ren, W., Heide, F., Nie, Y.: A survey on computational spectral reconstruction methods from rgb to hyperspectral imaging. *Scientific Reports* **12**(1), 11905 (Jul 2022). <https://doi.org/10.1038/s41598-022-16223-1>
33. Zhao, Y., Guo, H., Ma, Z., Cao, X., Yue, T., Hu, X.: Hyperspectral imaging with random printed mask. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10149–10157 (2019)
34. Zhou, Z., Zhang, Y., Xie, Y., Huang, T., Li, Z., Chen, P., Lu, Y.q., Yu, S., Zhang, S., Zheng, G.: Electrically tunable planar liquid-crystal singlets for simultaneous spectrometry and imaging. *Light: Science & Applications* **13**(1), 242 (Sep 2024). <https://doi.org/10.1038/s41377-024-01608-w>