



**HAL**  
open science

# The Marginal Importance of Distortions and Alignment in CASSI systems

Léo Paillet, Antoine Rouxel, Hervé Carfantan, Simon Lacroix, Antoine  
Monmayrant

► **To cite this version:**

Léo Paillet, Antoine Rouxel, Hervé Carfantan, Simon Lacroix, Antoine Monmayrant. The Marginal Importance of Distortions and Alignment in CASSI systems. 2025. hal-04903751

**HAL Id: hal-04903751**

**<https://laas.hal.science/hal-04903751v1>**

Preprint submitted on 21 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# The Marginal Importance of Distortions and Alignment in CASSI systems

Léo Paillet, Antoine Rouxel, Hervé Carfantan, Simon Lacroix and Antoine Monmayrant

**Abstract**—This paper introduces a differentiable ray-tracing-based model that incorporates aberrations and distortions to render realistic coded hyperspectral acquisitions using Coded-Aperture Spectral Snapshot Imagers (CASSI). CASSI systems can now be optimized in order to fulfill simultaneously several optical design constraints as well as processing constraints. Four comparable CASSI systems with varying degree of optical aberrations have been designed and modeled. The resulting rendered hyperspectral acquisitions from each of these systems are combined with five state-of-the-art hyperspectral cube reconstruction processes. These reconstruction processes encompass a mapping function created from each system’s propagation model to account for distortions and aberrations during the reconstruction process. Our analyses show that if properly modeled, the effects of geometric distortions of the system and misalignments of the dispersive elements have a marginal impact on the overall quality of the reconstructed hyperspectral data cubes. Therefore, relaxing traditional constraints on measurement conformity and fidelity to the scene enables the development of novel imaging instruments, guided by performance metrics applied to the design or the processing of acquisitions. By providing a complete framework for design, simulation and evaluation, this work contributes to the optimization and exploration of new CASSI systems, and more generally to the computational imaging community.

**Index Terms**—Compressive Sensing, Coded-Aperture, Hyperspectral Imaging, CASSI, Optical Design, Ray-tracing

## I. INTRODUCTION

**T**RADITIONAL hyperspectral imagers render three-dimensional data cubes by scanning the scenes along a spectral or spatial dimension [1]. This leads to two main drawbacks: the transmission and processing of a large quantity of data, and a limitation to static scenes. Snapshot hyperspectral imaging systems [2]–[6] rely on compressed sensing theory [7] to address these shortcomings. Specifically, CASSI systems [3], [6] reduce the redundancy present in hyperspectral scenes (HSSs) by performing a spatio-spectral encoding of the contained information. These systems use a coded aperture (or “mask”) and dispersive elements to spectrally and spatially modulate the HSS. The coded aperture can be optimized for specific applications [8], but can also be used as a random sampler [9], [10].

A primary goal of coded hyperspectral imaging is the reconstruction of the full hyperspectral cube (HSC), enabling further processing of the three-dimensional data cube. Historically, this reconstruction from coded acquisitions predominantly

relied on model-based methods [10]–[14]. While these methods offer insight into the reconstruction process, they suffer from long reconstruction times and suboptimal reconstruction quality. Recently, deep learning approaches have provided fast reconstruction times and high-quality results, albeit with less interpretability [8], [15]–[17]. Transformers [18] have also achieved superior results by leveraging non-local spatial relationships between pixels [19], [20], resulting in better exploitation of spatial and spectral interrelationships in HSSs.

However, state-of-the-art reconstruction algorithms [19], [21]–[24], while primarily data-driven, also integrate a simplified representation of the propagation model. They assume that CASSI systems produce spatially uniform dispersion, typically linear, without any optical misalignment. This is well suited for most CASSI systems that employ a double-Amici prism as the dispersive element [9], [25]–[29]. Compared to single prisms, the advantage of double-Amici prism assemblies is twofold: they allow for a direct-view geometry and do not exhibit anamorphic or optical distortions. Still, double-Amici prisms require longer manufacturing times, are more expensive, and misalignments must still be addressed when working with prototypes. Alternatively, using a single prism as the dispersive element results in a peculiar arrangement of information in the captured images, which can be properly exploited for reconstruction if the optical model is precise and accurate.

In this article, we evaluate the impact of distortions and misalignments on the coded information, when the optical model is properly considered. For this matter, we need a CASSI simulator accounting for both distortions and point spread function (PSF). Some image formation models [25] include PSF but they depend on calibration data, necessitating the assembly and calibration of a prototype. Recently, we proposed an accurate chief-ray-based propagation model [30] for dimensioning and optical distortions estimation, but it does not account for optical aberrations and cannot be used for Monte-Carlo rendering or PSF estimation. Recent works in computational imaging utilize differentiable ray-tracing to enable end-to-end optical design [31]–[33], in particular dO [33]. However, to the best of our knowledge, this approach has not yet been applied to coded aperture hyperspectral systems.

The most straightforward approach for us was to extent the dO rendering framework for hyperspectral systems. We propose implementing CASSI systems within dO, which enables accurate modeling of optical behavior and precise rendering of coded hyperspectral acquisitions. By leveraging ray tracing, we create an accurate spectral and spatial mapping between the object and image planes of CASSI systems. This mapping is then introduced in the reconstruction algorithms.

This work was supported by the CNRS and the Université de Toulouse. Léo Paillet is with LAAS-CNRS, Toulouse, France and with IRAP, Toulouse, France (e-mail: leo.paillet@laas.fr). Antoine Rouxel, Simon Lacroix and Antoine Monmayrant are with LAAS-CNRS, Toulouse, CNRS, France (e-mails: *firstname.name*@laas.fr). Hervé Carfantan is with IRAP, Toulouse, France (e-mail: Herve.Carfantan@irap.omp.eu).

This approach facilitates seamless processing across different CASSI systems, regardless of distortions or misalignments. We evaluate four distinct CASSI configurations, constructed with either a double-Amici prism assembly or a single prism. To assess the impact of misalignments, two of these configurations are deliberately misaligned. The reconstructions from these four configurations are compared using five reconstruction algorithms and standard evaluation metrics.

Our results demonstrate that distortions and misalignments have a marginal impact on the information encoded in CASSI acquisitions, and hence on the reconstruction quality, provided the rendering is realistic and an accurate model is incorporated into the reconstruction. This finding implies that the choice and fine-tuning of the reconstruction algorithm are more critical than the specific optical system.

The contributions of our work are:

- We implement coded aperture hyperspectral optical systems within a differentiable ray-tracing framework, enabling the rendering of distorted and aberrated images and facilitating the end-to-end design of such systems.
- We devise an accurate spatio-spectral mapping based on a realistic propagation model and introduce it, as a spatio-spectral prior, in the state-of-the-art reconstruction algorithms. This approach leverages two-dimensional measurements from the CASSI system and the associated rendering model.
- We highlight the marginal importance of distortions and misalignments in CASSI systems for acquisition processing purposes, provided the reconstruction process incorporates an accurate propagation model.

We first design a double-Amici prism assembly with the same angular spectral spreading than an off-the-shelf single prism, but none of its optical distortions. We then design two comparable CASSI systems –one based on the single prism, one on the double-Amici prism assembly–, and consider two configurations for each : with perfectly aligned and purposefully misaligned prisms. Those four configurations are designed in the implemented differentiable ray-tracing renderer outputting realistic CASSI acquisitions. Afterwards, we process the coded acquisitions to reconstruct HSCs with algorithms taking into account the propagation model. We finally compare the reconstruction quality reached with all four configurations, for five state-of-the-art algorithms.

All the source code about our modified version of dO, the adapted networks, and how to generate the figures is available at <https://github.com/lpaillet-laas/DiffCassiSim>.

## II. OPTICAL DESIGN AND RENDERING

Given that CASSI systems performances cannot be fairly assessed and compared using simple propagation models [10], [34], [35], which overlook the geometric distortions caused by dispersive elements as well as misalignment, we have developed a differentiable ray-tracing-based renderer based on dO, that takes into account distortions, misalignments and aberrations.

We applied it to generate accurate rendering of the coded acquisitions of two single disperser (SD) CASSI systems: one




Fig. 1: Top-view of the systems layout for the (SP) configuration (a) and the (AP) configuration (b).

system using a single prism as the dispersive element, referred as (SP), and one with a double-Amici prism assembly, referred as (AP), shown in Figure 1. The two other configurations are duplicates for which the dispersive element has been purposefully misaligned by  $5^\circ$  around the  $x$ -axis with respect to the reference frames of Figure 1, and respectively referred as (mSP) and (mAP).

In this section, we delve into:

- The design process of a custom double-Amici prism to minimize optical distortions and achieve direct-view.
- The implementation of the four optical configurations within dO, highlighting how different configurations are simulated.
- The methodology used for generating realistic coded hyperspectral acquisitions.

### A. Case Study: SD-CASSI

The (SP) system illustrated in Figure 1-(a) contains a first lens to collimate the light onto the prism and a second one to image the light onto the detector. For the (AP) system, as for most recent Amici-based systems [9], [26], a single relay lens is used in a  $2f - 2f$  configuration, with the dispersive element positioned between the lens and the detector, as shown in Figure 1-(b).

To ensure a fair comparison between both systems, it is crucial to dimension the dispersive elements so that they exhibit comparable spectral dispersion. Our analysis is conducted on systems with a detector comprising  $512 \times 512$  pixels with a  $10\text{-}\mu\text{m}$  pitch resulting in a  $\simeq 5 \times 5 \text{ mm}^2$  field of view, a spectral range of  $[450 - 650]$  nm centered at 520 nm, and a numerical aperture of 0.05. The systems utilize lenses with a focal length of  $f = 50 \text{ mm}$ . The targeted angular spectral spreading for these systems is  $\Delta_0 = 0.95^\circ$ , resulting in a spatio-spectral spreading  $S$  on the detector of  $S \simeq 830 \mu\text{m}$  (83 pixels) for the central point of the field of view. These parameters are selected based on the following criteria: the expected number of resolved pixels across the field of view, the consistency with the spectral range and resolution of the KAIST dataset [36], and the compatibility with standard off-the-shelf components.




Fig. 2: Optical designs of dispersive elements in our imaging system. (a) A commercially available N-BK7 equilateral prism, aligned in for minimum deviation  $D_0$ . (b) A custom-designed double-Amici prism optimized to minimize distortions and achieve direct-view geometry. Both designs exhibit comparable angular spectral spreading  $\Delta$  to ensure a fair comparison of system performance.

### B. Design of a Double-Amici Prism for Fair Comparison

Our first goal is to design a double-Amici prism that replicates the angular spectral dispersion  $\Delta_0 = 0.95^\circ$  of a standard N-BK7 equilateral prism at minimum deviation  $D_0$  illustrated in Figure 2-(a), while minimizing geometric distortions and ensuring a direct-view configuration.

To design the prism, we employ the chief-ray-based dimensioning tool SIMCA, described in [37], [38], which solves a gradient-based optimization problem incorporating both optical and system-related parameters. Although dO is more powerful for differentiable optical design, we chose SIMCA as it was specifically developed for CASSI systems, thus offering more straightforward and faster optimization.

1) *Prism Parametrization*: We optimize the following five parameters, illustrated in Figure 2-(b):  $\alpha_c$  the angle of incidence on the double-Amici prism assembly,  $A_1$  and  $A_2$  the apex angles of the first and second prisms, respectively, and  $m_1$  and  $m_2$  the glass materials for the two prisms, selected from the Schott catalog. Since the glass materials are discrete and not differentiable by default, we model the dispersion curve of each glass using two continuous sub-parameters: the refractive index at the "d" Fraunhofer line and the Abbe number, following the methodology described in [31], [39]. During optimization, we treat these sub-parameters as continuous variables. After optimization, we select the glass materials from the catalog that are closest to the optimized sub-parameter values.

2) *Loss Functions*: To optimize the prism design, we first aim to match the angular spatio-spectral dispersion of the base design  $\Delta_0$ . We then minimize the optical distortions across the field of view and for the whole spectral range. We finally enforce an easy-to-align and compact system by minimizing the angular deviation  $D$  and the thickness of the prism.

The *spectral dispersion loss*  $\mathcal{L}_\Delta$  is defined as the squared difference between the base spectral dispersion  $\Delta_0$  and the dispersion  $\Delta$  calculated for the current prism design. We compute  $\Delta$  as the absolute difference between the output angles after the prism for the shortest and longest wavelengths at the center of the field-of-view, respectively denoted as  $\Delta_{\lambda_{min}}$  and  $\Delta_{\lambda_{max}}$ :

$$\mathcal{L}_\Delta = (\Delta_0 - |\Delta_{\lambda_{max}} - \Delta_{\lambda_{min}}|)^2 \quad (1)$$

The *distortion loss*  $\mathcal{L}_\varepsilon$  quantifies the geometric distortions introduced by the prisms. We calculate a distance tensor  $\varepsilon$  which measures how each imaged point of the scene is displaced due to distortion. In practice, we measure the Euclidean distance between the distorted coordinates  $(X_{ds}, Y_{ds})$  and the ideal coordinates  $(X_{id}, Y_{id})$ :

$$\varepsilon_{i,j,k} = \sqrt{(X_{ds}^{i,j,k} - X_{id}^{i,j,k})^2 + (Y_{ds}^{i,j,k} - Y_{id}^{i,j,k})^2} \quad (2)$$

Here,  $(X_{id}^{i,j,k}, Y_{id}^{i,j,k})$  are the coordinates of the ideal (undistorted) image grid points as described in [30]. We then define the distortion loss as the square of the maximum value in this distance tensor:

$$\mathcal{L}_\varepsilon = \left( \max_{i,j,k} \varepsilon_{i,j,k} \right)^2 \quad (3)$$

The minimization of  $\mathcal{L}_\varepsilon$  reduces the maximum geometric distortion across all points in the image grid.

The *deviation loss*  $\mathcal{L}_D$  accounts for the total angular deviation induced by the prism configuration. The deviation is computed based on the chief-ray angles  $\alpha_c$  and  $\alpha_c^{out}$  (incidence angle and output angle of the chief-ray at the central wavelength) and the apex angles of the prisms. The deviation loss is then given by the squared total deviation:

$$\mathcal{L}_D = (\alpha_c + \alpha_c^{out} + 2A_1 - A_2)^2 \quad (4)$$

where  $A_i$  are the apex angles of the prisms.

The *thickness loss*  $\mathcal{L}_t$  approximates the physical thickness of the double-Amici prism. It is proportional to the sum of the squared apex angles:

$$\mathcal{L}_t = 2A_1^2 + A_2^2 \quad (5)$$

This loss helps to minimize the overall size and weight of the optical system.

The *glass distance loss* ( $\mathcal{L}_g$ ) measures the squared difference between the refractive index ( $n_d$ ) and Abbe number ( $v_d$ ) of the selected glass materials and the closest available materials in the Schott catalog. This loss ensures that the materials chosen during the optimization process are realistic and match closely with available catalog materials. For a double-prism system, the glass distance loss is defined as:

$$\mathcal{L}_g = \min_{\text{glass}} \left( \left( \frac{n_d - n_d^{\text{catalog}}}{\Delta n_d} \right)^2 + \left( \frac{v_d - v_d^{\text{catalog}}}{\Delta v_d} \right)^2 \right) \quad (6)$$

where  $\Delta n_d$  and  $\Delta v_d$  represent the ranges of refractive index and Abbe number in the Schott catalog. This ensures that the chosen materials are practical and manufacturable.

Finally, the *total internal reflection loss* ( $\mathcal{L}_R$ ) penalizes designs that approach the critical angle for total internal reflection, preventing undesirable optical properties such as significant losses or distortions due to reflection instead of the expected transmission. It is computed as:

$$\mathcal{L}_R = \text{Softplus}(2 \cdot \min(D_{\text{TIR}}))^2 \quad (7)$$

where  $D_{\text{TIR}}$  is the distance from which total internal reflection would occur. The Softplus function ensures that the loss only




Fig. 3: Distortion maps of a regular grid of points traced through the four considered optical systems, for the extrema and center wavelengths (450 nm, 520 nm, and 650 nm). Note the different colorscales between the configurations.

increases significantly when the angle approaches the critical threshold, preventing abrupt changes in the optimization process and ensuring smooth convergence.

3) *Design Optimization*: The goal is to find the set of design parameters  $\theta = \{\alpha_c, A_1, A_2, m_1, m_2\}$  that minimize the total loss  $\mathcal{L}(\theta)$  defined as a linear combination of the 6 loss terms:

$$\mathcal{L} = w_{\Delta}\mathcal{L}_{\Delta} + w_{\varepsilon}\mathcal{L}_{\varepsilon} + w_D\mathcal{L}_D + w_t\mathcal{L}_t + w_g\mathcal{L}_g + w_R\mathcal{L}_R \quad (8)$$

$\mathcal{L}$  encapsulates both optical and system-related objectives, allowing for a comprehensive optimization of the prism design. The weights  $w$  in the loss expression 8 balance the contributions of each loss term according to their importance in the design objectives, they have been empirically chosen to guide the optimization (the values used in the optimization are  $(w_{\Delta}, w_{\varepsilon}, w_D, w_t, w_g, w_R) = (1, 1, 2.5 \times 10^6, 5 \times 10^3, 10^{10} \times \text{iteration\_number}, 10)$ ).

4) *Results*: Using the Adam optimizer [40] and starting from the Amici parameters described in [9], we design a double-Amici prism assembly that meets our spectral dispersion requirement  $\Delta_0$  of  $0.95^\circ$  at the detector plane within  $\pm 1\%$ . The final materials for the Amici prism are N-SK2 and SK10 for  $m_1$  and  $m_2$ , respectively. The apex angles are  $A_1 = 29.2^\circ$  and  $A_2 = 47.9^\circ$ , with the incident chief-ray arriving at the prism at  $\alpha_c = 5.1^\circ$ .

With respect to the (SP) setup, the maximum geometric distortion is reduced from  $214\mu\text{m}$  to  $6\mu\text{m}$  and the mean distortion from  $75\mu\text{m}$  to  $1.8\mu\text{m}$  over a  $5 \times 5\text{ mm}$  field of view. Figure 3 shows the distortion map for both systems (SP) and (AP), and also for the misaligned configurations (mSP)

and (mAP). This distortion minimization is achieved while maintaining a direct-view geometry with a total deviation  $D$  of  $0.1\text{ mrad}$ .

Both the commercially available N-BK7 equilateral prism and our custom-designed double-Amici prism assembly exhibit comparable spectral dispersion while presenting different geometric distortions, ensuring a fair comparison of systems performances in section III-D.

### C. Differentiable Simulation for SD-CASSI systems

Both single prism and double Amici prism assembly are employed to implement SD-CASSI systems using the differentiable optics design tool dO.

Figure 1 shows the two system's configurations in dO. As stated in section II-A, (AP) and (mAP) use a single relay lens (see Figure 1-(b)). Typically, an achromatic objective lens is utilized to ensure high spatial resolution across the field of view and the entire spectral range. However, in our implementation (AP), the objective lens is modeled as an ideal thin lens to provide a generalized analysis that focuses on prism-related aberrations and geometric distortions, irrespective of the objective lens used. Additionally, replacing an objective by an ideal thin lens has no significant impact on the simulations when considering our pixel sizes and field of view. In the single-prism-based configurations (SP) and (mSP), two lenses are on either side of the prism (see section II-A). As proposed in [41], it is advantageous to minimize aberrations early in the setup by employing an achromatic objective lens before the prism and a simple doublet lens afterwards to reduce costs. Similar to the approach used for the Amici system, we model the achromatic objective lens as an ideal thin lens and implement a Thorlabs AC254-050-A-ML doublet lens for the second lens, as illustrated in Figure 1-(a).

The misaligned configurations (mSP) and (mAP) have been simulated by rotating the dispersive element by  $5^\circ$  around the  $x$ -axis as stated in Section II.

1) *Differentiable Ray-Traced Simulations*: dO supports both forward and backward ray-tracing, enabling the optimization of the design of imaging systems and the rendering of scenes through these systems. It was originally developed for designing free-form optics RGB imaging systems with axial symmetry: adapting it to CASSI systems required modifications due to the loss of axial symmetry caused by the prism. These modifications included the addition of a prism optical element, based on dO built-in surfaces, and the ability to perform smooth rotations between optical elements, ensuring accurate modeling and simulation of the CASSI architecture.

In our modified dO, each optical element is then treated as an independent system, and rays are traced between these systems. The optical elements are composed of surfaces separated by materials, as in Zemax. Each element can then be rotated and shifted individually to closely match the design specifications.

The implementation of SD-CASSI designs in dO enables the acquisition of hyperspectral and coded hyperspectral images that realistically reflect the given optical system, in contrast to the simplified mathematical models used in [3], [11], [35].



Such models often struggle to account for all distortions and dispersions present in real optical systems. Our approach is more aligned with the mathematical models of [25], [42]. In addition, it directly generates PSFs through accurate ray tracing and can account for spectral dispersions along the  $y$  (vertical) axis and for the spectral dispersions continuity thanks to Monte-Carlo methods. This is particularly evident in the case of misaligned optical configurations, as seen in the (mAP) and (mSP) configurations (refer to Figure 3 and Supplementary Figure 13) although the Amici-based system is more tolerant to misalignment.

Additionally, dO supports the differentiable optimization of these designs, allowing fine-tuning of rotations, shifts between optical elements, and adjustments of surface distances and configurations. This functionality allows both the design and operational use of SD-CASSI systems within the same application, facilitating seamless imaging co-design.

2) *Validation of the Implementation in dO*: To validate our dO implementation, we compared the PSFs from dO and Zemax of all four configurations across various wavelengths and spatial positions. Figure 4 shows a selection of four PSFs for various positions across the field of view for the (mSP) configuration at a wavelength of  $\lambda = 650$  nm (more PSFs are presented in the supplementary material, Figures 8 to 11). Although this configuration exhibits the strongest distortions and is thus the most challenging of the four considered configurations, the PSFs modelled with dO are similar to their Zemax counterparts. The spot diagrams from both implementations (dO and Zemax) are highly consistent, with Root-Mean-Square (RMS) spot size differences below  $1.5 \mu\text{m}$ , well below the considered  $10\text{-}\mu\text{m}$  pixel size, thereby ensuring adequate rendering resolution.

Differences observed in spot diagrams primarily arise from aperture handling which impacts ray distribution: dO assumes a telecentric system without an aperture stop, whereas Zemax includes an aperture object. Adding an aperture function to dO would require ray-direction calculations for aperture filling, but the benefits would be minimal given the minor discrepancies between both systems. Additionally, ray sampling methods differ: with dO we employ hexapolar sampling, while Zemax uses grid sampling for PSF determination.

Further validation was conducted by comparing distortions at various wavelengths and positions obtained with both dO and Zemax. Figure 5 shows the difference between Zemax and dO distortions across the field of view for extrema and central wavelengths. Results show that distortions are accurately modeled, with mean differences of only  $1 \mu\text{m}$  across the three wavelengths, a tenth of the actual pixel size. Minor discrepancies in distortion maps between dO and Zemax also stem from their differing aperture treatments.

#### D. Coded acquisition rendering

1) *Rendering process*: Our model of the four optical configurations enables HSS rendering through backward ray-tracing. Each spectral plane is rendered individually, and acquisition is performed by summing along the spectral dimension. We render scenes with 28 spectral bands evenly distributed from

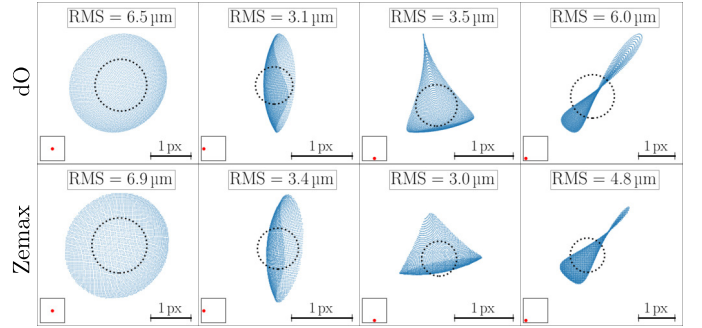


Fig. 4: PSFs obtained with the misaligned single prism configuration (mSP) at four positions in the field of view (positions denoted by the red points on the bottom left of each figure). Top: PSFs obtained with dO, bottom: PSFs obtained with Zemax. The dotted black circle corresponds to the RMS radius centered on the centroid of the PSFs.




Fig. 5: Maps of the difference between distortions computed with dO and Zemax, for the extrema and central wavelengths (450 nm, 520 nm, and 650 nm). Note the same colorscales across the 4 configurations.

450 nm to 650 nm from the hyperspectral datasets CAVE [43] and KAIST [36]. Due to the systems'  $830 \mu\text{m}$  nonlinear spectral spread on the detector and  $10 \mu\text{m}$  spatial sampling, the data is oversampled to ensure accurate rendering over the 83 illuminated pixels (otherwise some pixels would not receive any signal, giving a spatially non-continuous rendered acquisition). The smallest integer oversampling factor  $n$  so that  $28 \times n \geq 83$  is  $n = 3$ : we oversample with  $n = 4$  to alleviate missing signal errors that could occur due to Monte-Carlo sampling with a small amount of rays.

Considered HSSs therefore contain  $4 \times 28 = 112$  spectral

bands. For a fast rendering of the  $512 \times 512 \times 112$  hyperspectral scenes, we limit ray tracing to 20 rays per pixel per wavelength. Following the modifications done to dO, rendering is performed sequentially, progressing backwards through each optical element and independently for each wavelength.

To further mitigate quantization errors from Monte-Carlo sampling arising from tracing a small amount of rays, the rendered planes are convolved with a smoothing kernel. We convolve each spectral plane with an Airy disk, varying according to  $\lambda$ , whose diameter equals 2.5 pixels ( $= 25 \mu\text{m}$ ) at 520 nm, adequately smoothing the rendering through convolution.

To simulate a HSS acquisition  $\mathbf{H}$  with a given 2D binary mask  $\mathbf{M}$  through a SD-CASSI system, we compute the coded scene  $\mathbf{H}^c$  as follows:

$$\mathbf{H}^c(:, :, n_\lambda) = \mathbf{H}(:, :, n_\lambda) \odot \mathbf{M}, \forall n_\lambda \in [0, 112]$$

$\mathbf{H}^c$  is then input into our rendering process, representing information at the entry of the optical system.

2) *Validation of the Rendering Method:* We validate our rendering method with a single-slit mask, for which the SD-CASSI systems emulate an imaging prism spectrometer. With a single 1-pixel-wide slit opened on the mask  $\mathbf{M}$ , both aligned configurations (AP) and (SP) disperse the spectrum of the imaged slit along the  $x$  axis. Since distortions mainly appear in the (SP) configuration, we will only consider this harder case.

The test scene selected from the CAVE dataset is shown in Figure 6-(a), together with the chosen slit position (dashed line). The slit is located at the center of the field of view, and 3 regions with a constant spectrum are acquired (green, orange and gray squares in Figure 6-(a)). These spectra are spatially dispersed based on the system’s spectral spread, and convolved with the corresponding spatio-spectral PSFs and Airy disks. To account for the  $y$ -axis spread of the PSFs, we average 40 rows of the acquisition in each region and compare this result to the corresponding ground truth spectrum under the same experimental conditions specified earlier (see section II-D1): 20 rays per pixel per wavelength across 112 wavelengths. We upsample the ground truth spectra to 280 wavelengths to account for their continuity and to reduce the quantization errors that would otherwise appear in the comparison, with no physical meaning.

Results presented in Figure 6-(c) show a close match between the acquisitions (solid lines) and the ground truth spectra (dashed lines). Thus, our model effectively simulates a SD-CASSI system, enabling realistic acquisition simulations that align with those from a physical system.

We created a fork of dO including all our modifications that can be accessed at <https://github.com/lpaillet-laas/DiffOptics>.

Additionally, another repository accessible at <https://github.com/lpaillet-laas/DiffCassiSim>, contains this fork of dO together with all the code for the processing done in Section III.

### III. HYPERSPECTRAL CUBE RECONSTRUCTIONS

#### A. Overview

We demonstrate here that for the four considered configurations, a comparable amount of information is encoded

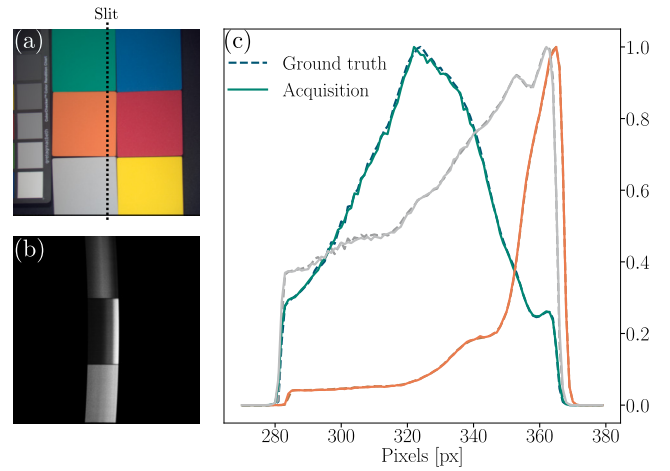


Fig. 6: Spectra extracted from rendered single-slit acquisition with the (SP) configuration compared to ground truth spectra retrieved from the HSS. The ground truth spectra correspond to the HSS data spatially dispersed and convolved by the PSFs and Airy disks of the system. (a) RGB image of the scene and location of the slit (dashed line). (b) Rendered acquisition with the (SP) configuration. (c) Acquired spectra (solid lines) in the 3 regions with constant color (green, orange and gray) compared to their ground truth counterparts (dashed lines).

in the acquisitions and that it can be retrieved by state-of-the-art approaches when the system model is known and considered, hence exhibiting the marginal impact of distortions and misalignments.

For this purpose, the hyperspectral cube reconstruction from coded acquisition is an ideal test case. Deep learning methods [8], [36], [44] being now the best solution for HSC reconstruction, we apply five recent state-of-the-art algorithms, and assess the reconstruction quality with a series of usual metrics. The considered reconstruction algorithms are Deep Gaussian Scale Mixture Prior (DGSMP) [21], Mask-guided Spectral-wise Transformer (MST) [19], Degradation-Aware Unfolding Half-Shuffle Transformer (DAUHST) [22], Residual Degradation Learning Unfolding Framework (RDLUF) [23], and Pixel Adaptive Deep Unfolding Transformer (PADUT) [24].

Our evaluation uses the two public HSS datasets CAVE [43] and KAIST [36]. The CAVE dataset includes 32 HSSs of size  $512 \times 512 \times 31$ , while the KAIST dataset includes 30 HSSs of size  $2704 \times 3376 \times 31$ . HSSs have been spectrally interpolated following [43] in order to be consistent with the wavelengths of the system used to acquire these HSSs. They thus contain 28 spectral bands ranging from 450 nm to 650 nm. To fit the field of view ( $\approx 5 \times 5 \text{ mm}^2$ ) with a pixel size of  $10 \mu\text{m}$ , each HSS in the datasets was cropped to reach a final size of  $512 \times 512 \times 28$ . Following the setups of the referenced networks, models were trained on the CAVE dataset and tested on 10 scenes taken from the KAIST dataset.

#### B. Spatio-Spectral Mapping and Reconstruction Initialization

Most networks leverage *a priori* information from the acquisition to reconstruct the hyperspectral cube. For perfectly

aligned systems without distortions, a scene  $\mathbf{H}$  with dimensions  $H \times W \times N_\lambda$  is rendered onto an acquisition  $\mathbf{A}$  of size  $H \times (W + S)$ , where  $S$  denotes the spectral spread (in pixels) of the optical system. Defining the spectral spread at a specific wavelength  $n_\lambda$  as  $s(n_\lambda)$ , the typical network initialization  $\mathbf{I}$  is given by:

$$\mathbf{I}(:, :, n_\lambda) = \mathbf{A}(:, s(n_\lambda) : W + s(n_\lambda)), \forall n_\lambda \in \llbracket 0, N_\lambda \rrbracket \quad (9)$$

resulting in an initialized cube of dimensions  $H \times W \times N_\lambda$ .

With negligible distortions, all information from the scene at wavelength  $n_\lambda$  is indeed captured in a rectangular area of size  $H \times W$  starting at the position  $x = s(n_\lambda)$ . However, as shown in Figure 3 and in Figure 13 of the supplementary material, some of our configurations exhibit spatial distortions, so information from each wavelength is not represented in a rectangular area. In the misaligned configurations (mAP) and (mSP), the acquisition  $\mathbf{A}$  also displays dispersion along the  $y$  axis, resulting in an acquisition of dimensions  $(H + S_x) \times (W + S_y)$ , where  $S_x$  denotes the spectral spread (in pixels) along the  $x$  axis, and  $S_y$  the spectral spread along the  $y$  axis.

Knowing both the configurations and our model, we can define a mapping  $f$  for each configuration such that:

$$f(x_s, y_s, n_\lambda) = (x_d, y_d) \quad (10)$$

where  $(x_s, y_s, n_\lambda)$  denotes a position in the scene and  $(x_d, y_d)$  is the corresponding position on the detector. This mapping  $f$  is generated by identifying the *pixel* locations of rays traced from a grid of points of dimensions  $H \times W$  for each of the  $N_\lambda$  wavelengths.

Using the mapping  $f$ , for each wavelength  $n_\lambda$ , we can accurately fit parts of the distorted acquisition into a rectangle of shape  $H \times W$ , thus initializing  $\mathbf{I}$  as:

$$\mathbf{I} = \mathbf{A} \circ f \quad (11)$$

This initialization ensures accurate computation, aligning with the model and the specific optical system in use.

The considered reconstruction algorithms contain a simplified propagation model to achieve better results than if they were purely data-driven. Therefore, these algorithms have been modified to incorporate the mapping function  $f$  in the reconstruction process, as it allows the computation of both the direct and adjoint operators used in the unfolded iterative steps. This allows an accurate propagation model to be integrated in the reconstruction algorithm, faithful to how acquisitions are generated.

### C. Networks Implementation

The networks were adapted to accommodate our rendered acquisitions. Originally designed to reconstruct  $256 \times 256 \times 28$  HSCs, the networks were modified to reconstruct  $512 \times 512 \times 28$  HSCs required by our configurations. Additionally, most networks were previously hardcoded for simplified optical setups assuming linear spectral dispersion with a slope of 2 pixels ( $20 \mu\text{m}$  in our case) across  $560 \mu\text{m}$ . However, our configurations exhibit unique non-linear spectral spread across

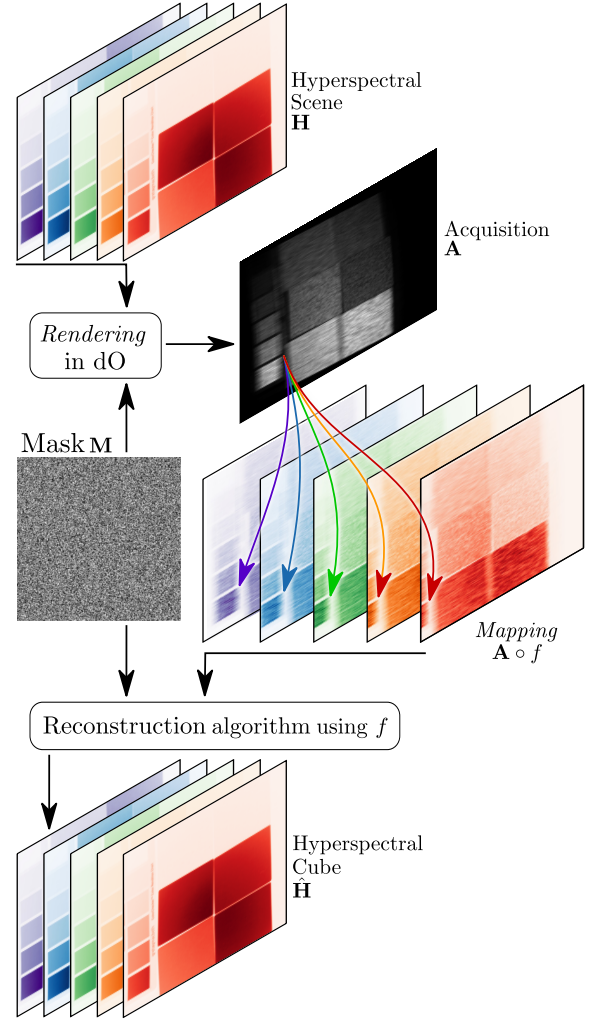


Fig. 7: Workflow to reconstruct HSCs with a given optical system and a given reconstruction algorithm. The scene  $\mathbf{H}$  is acquired with the optical system thanks to our rendering with dO. The acquisition is then mapped to an initialization using the model through  $f$ . This serves as an input to a reconstruction algorithm yielding  $\hat{\mathbf{H}}$ . The mapping function is also used in the algorithm.

$830 \mu\text{m}$ . The networks were therefore adapted with the mapping function  $f$  to process the acquisitions, as illustrated in Figure 7.

For all HSSs, a unique binary mask with a random opening ratio of 0.5 was used. Consistent with original methods, we use the Adam optimizer and both learning rates and scheduler schemes were maintained. Each network and configuration was trained for 400 epochs, utilizing random cropping, flipping, and rotation for data augmentation. For the (mAP) configuration, DGSMP was trained with a learning rate of  $0.8 \times 10^{-4}$  and 500 epochs since the training process was not steady with a larger learning rate of  $10^{-4}$ . The batch size was set to 1, with gradient accumulation every 4 batches, and training conducted on NVIDIA RTX A100 and A6000 GPUs.

Training loss was calculated as the Root Mean Square Error (RMSE) between the ground truth HSSs and reconstructed



HSCs. Evaluation metrics include RMSE, Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Spectral Angle Mapper (SAM) [45]. The SAM metric allows to evaluate the spectral accuracy of the reconstructions.

#### D. Results

		(AP)	(SP)	(mAP)	(mSP)
<b>RMSE</b> ↓ ( $\times 10^{-3}$ )	DGSMP	36.7	39.1	<b>35.4</b>	40.7
	MST	27.4	<b>26.6</b>	<b>26.6</b>	28.5
	DAUHST	22.1	<b>20.0</b>	20.3	20.7
	RDLUF	<b>23.0</b>	25.3	24.4	25.4
	PADUT	21.8	<b>19.8</b>	22.6	22.1
<b>PSNR</b> ↑	DGSMP	27.1	26.6	<b>27.4</b>	26.3
	MST	29.6	<b>29.9</b>	<b>29.9</b>	29.3
	DAUHST	31.4	<b>32.3</b>	32.2	32.1
	RDLUF	<b>31.1</b>	30.2	30.6	30.3
	PADUT	31.6	<b>32.4</b>	31.2	31.5
<b>SSIM</b> ↑ [0 – 1]	DGSMP	<b>0.855</b>	0.826	0.829	0.843
	MST	0.909	<b>0.913</b>	0.910	0.900
	DAUHST	0.933	<b>0.943</b>	0.939	0.937
	RDLUF	<b>0.928</b>	0.918	0.921	0.917
	PADUT	0.934	<b>0.942</b>	0.931	0.933
<b>SAM</b> ↓ [0 – 1]	DGSMP	<b>0.074</b>	0.092	0.082	0.083
	MST	0.057	<b>0.053</b>	0.057	0.058
	DAUHST	0.054	<b>0.051</b>	<b>0.051</b>	0.056
	RDLUF	<b>0.054</b>	0.056	0.057	0.055
	PADUT	0.051	<b>0.049</b>	0.053	0.052

TABLE I: Average RMSE, PSNR (in dB), SSIM and SAM on all test scenes, for each optical system and reconstruction algorithm.

As stated earlier, the networks were trained with our four configurations on the CAVE dataset and were then tested on the KAIST dataset. The quantitative results are presented in Table I. SAM has been normalized between 0 and 1 to be in the same range as SSIM. We compare the evaluation metrics across all four configurations for each state-of-the-art reconstruction algorithm. For some networks, the evaluation metrics are slightly lower than those reported in corresponding papers. This occurs because we adapted the networks to process  $512 \times 512 \times 28$  HSCs, although they were originally designed to reconstruct  $256 \times 256 \times 28$  HSCs. Thus, achieving the same performance might require either more training epochs or a deeper architecture to better capture small spatial variations.

As seen in Table I, for a given reconstruction algorithm, all four evaluation metrics show no significant difference across all four configurations. The largest difference between configurations is observed with DGSMP, where (AP) yields a SSIM of 0.855 and (SP) yields a SSIM of 0.826, representing a 0.029 SSIM difference. For other reconstruction algorithms, the SSIM difference does not exceed 0.013. For further details, the metrics for each scene can be seen in the supplementary materials (see Supplementary Tables IV to VIII). These non-significant differences indicate that the encoded information quality remains consistent across all four configurations, despite distortions and misalignments in some

System	Training	Test	PSNR ↑	SSIM ↑
(SP)	Simple	Simple	31.1	0.928
	Simple	Real	19.9	0.736
	Real	Real	<b>32.4</b>	<b>0.942</b>
(mSP)	Simple	Simple	<b>31.9</b>	<b>0.934</b>
	Simple	Real	15.1	0.571
	Real	Real	31.5	0.933

TABLE II: Rendering ablation result, with PADUT algorithm.

System	Mapping	PSNR ↑	SSIM ↑
(SP)	Without	31.1	0.932
	With	<b>32.4</b>	<b>0.942</b>
(mSP)	Without	29.8	0.914
	With	<b>31.5</b>	<b>0.933</b>

TABLE III: Mapping ablation result, with PADUT algorithm.

configurations. Moreover, the configuration yielding the best evaluation metrics varies with the algorithm used. This further validates that processing results mainly depend on the quality of the processing algorithm, rather than the optical system, provided the system is accurately modeled. Thus, we can relax certain constraints on optical design while achieving the same processing performance.

#### E. Ablation Study

An ablation study was conducted to evaluate the impact of rendering accuracy and mapping on HSC reconstruction. The study was performed on two distortion-inducing configurations: (SP) and (mSP), and exclusively employed the PADUT algorithm, as it provided the highest reconstruction quality.

In the first experiment, we assessed the effect of incorrect rendering during the training phase. PADUT was trained using a simplified rendering, and reconstruction quality metrics were evaluated with our realistic rendering. The simplified model renders spectral planes using the (AP) system, which introduces negligible distortions and serves as the baseline for standard simplified frameworks [9], [11], [35]. To avoid bias in reconstruction, the spatio-spectral spreading  $S$  curve was set to match that of the realistic propagation model, minimizing minor differences between the systems (see Supplementary Figure 12). Spatial shifts were thus applied to the rendered spectral planes to account for spatio-spectral spreading mismatches before summing them to form the acquisitions.

In the second experiment, the mapping function  $f$  was entirely removed from the workflow, meaning a simplified propagation model is used. Indeed, only the horizontal spatio-spectral spreading characteristics of the configurations were used for initial reconstruction and during the algorithm’s processing, to account for the correct dispersion and to prevent bias. This experiment isolates the influence of the mapping function on reconstruction quality. Both (SP) and (mSP) configurations were considered to account for distortions and misalignments, with the mapping excluded during training and testing.

Results for the first experiment are presented in Table II. They demonstrate that using different rendering models be-

tween the training and test phase significantly degrades reconstruction quality. However, using the same rendering method for both phases gives similar performances for all configurations, as seen in section III-D. This highlights the fact that reconstruction algorithms cannot be effectively trained on acquisitions rendered using simplified models with the expectation of robust performance when tested on systems employing realistic, non-simplified rendering processes, or on real acquisitions from CASSI prototypes.

In the second experiment (see Table III), the removal of the mapping function  $f$  from the reconstruction process only caused a slight deterioration in reconstruction quality. We hypothesize that the limited impact is due to the narrow spatio-spectral spread difference between consecutive wavelengths, which allows missing spatial information to be interpolated from neighboring wavelengths during training. Still, excluding a correct mapping during the training phase notably reduces the achievable quality metrics, thereby amplifying the performance gap relative to systems without distortions or misalignments.

Ultimately, rendering accuracy during both training and test has a much larger influence on reconstruction quality than the inclusion of a correct mapping function. However, a proper mapping still contributes to improved algorithm performance, providing a non-negligible enhancement in reconstruction accuracy. Both steps are therefore important to reach the best reconstruction quality with a given algorithm.

#### IV. CONCLUSION

We focused on the design and performance evaluation of Coded Aperture Snapshot Spectral Imaging (CASSI) systems. Our work aimed to bridge the gap between optical hardware and computational processing by providing realistic simulations and analyses that consider the complexities of actual optical setups.

Our first contribution was the implementation of coded aperture hyperspectral optical systems within a differentiable ray-tracing framework. This enables to render synthetic coded hyperspectral images that accurately incorporate optical distortions and aberrations. By leveraging this framework, we could simulate four optical designs more precisely than with simplified mathematical models. Secondly, a realistic propagation model was utilized to map 2D coded measurements to estimations of hyperspectral cubes, serving both as the initialization and throughout the reconstruction process. This approach allowed for improved reconstruction by accounting for the exact optical characteristics of the configuration, including geometric distortions and spectral dispersion. This workflow can be generalized to a great variety of systems with our framework, given the propagation model is known. Thirdly, we demonstrated that geometric distortions and misalignments in CASSI systems have a marginal impact on reconstruction performance. Our evaluations showed that the choice of the reconstruction algorithm plays a more critical role in determining the quality of the reconstructed hyperspectral cubes than the specific optical system used, provided the system is accurately modeled, incorporated into the reconstruction

process and used to render realistic acquisitions. We conclude that the same amount of information is transmitted regardless of the system used.

Looking forward, this work yields the comparison of different CASSI system performances using information theory measures, facilitating more informed design choices [46]–[48], with no need for a processing algorithm to evaluate the design. It additionally opens several avenues for future research. The differentiability of the ray-tracing simulator can be exploited for optimizing coded apertures, potentially leading to designs that maximize information capture or minimize reconstruction error. Dynamic mask designs could also be achieved, accurately adapting to a scene in order to process several coded acquisitions and reach better information acquisition. Finally, this work paves the way for end-to-end optimization of CASSI systems, integrating both optical design and computational algorithms to achieve optimal performance in co-design computational imaging.

#### ACKNOWLEDGMENTS

This project has received financial support from the CNRS through the MITI interdisciplinary programs.

#### REFERENCES

- [1] Y. Garini, I. T. Young, and G. McNamara, "Spectral imaging: principles and applications," *Cytometry A*, vol. 69, no. 8, pp. 735–747, Aug. 2006.
- [2] X. Yuan, D. J. Brady, and A. K. Katsaggelos, "Snapshot compressive imaging: Theory, algorithms, and applications," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 65–88, 2021.
- [3] A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Appl. Opt.*, vol. 47, no. 10, pp. B44–B51, Apr 2008. [Online]. Available: <https://opg.optica.org/ao/abstract.cfm?URI=ao-47-10-B44>
- [4] G. R. Arce, D. J. Brady, L. Carin, H. Arguello, and D. S. Kittle, "Compressive coded aperture spectral imaging: An introduction," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 105–115, 2014.
- [5] X. Cao, T. Yue, X. Lin, S. Lin, X. Yuan, Q. Dai, L. Carin, and D. J. Brady, "Computational snapshot multispectral cameras: Toward dynamic capture of the spectral world," *IEEE Signal Processing Magazine*, vol. 33, no. 5, pp. 95–108, 2016.
- [6] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express*, vol. 15, no. 21, pp. 14 013–14 027, Oct 2007. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-15-21-14013>
- [7] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [8] L. Wang, T. Zhang, Y. Fu, and H. Huang, "HyperReconNet: Joint coded aperture optimization and image reconstruction for compressive hyperspectral imaging," *IEEE Trans. Image Process.*, vol. 28, no. 5, pp. 2257–2270, Nov. 2018.
- [9] A. Wagadarikar, N. P. Pitsianis, X. Sun, and D. J. Brady, "Video rate spectral imaging using a coded aperture snapshot spectral imager," *Opt. Express*, vol. 17, no. 8, pp. 6368–6388, Apr 2009. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-17-8-6368>
- [10] I. Ardi, H. Carfantan, S. Lacroix, and A. Monmayrant, "Fast hyperspectral cube reconstruction for a double disperser imager," in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, sep 2018.
- [11] D. Kittle, K. Choi, A. Wagadarikar, and D. J. Brady, "Multiframe image estimation for coded aperture snapshot spectral imagers," *Appl. Opt.*, vol. 49, no. 36, pp. 6824–6833, Dec 2010. [Online]. Available: <https://opg.optica.org/ao/abstract.cfm?URI=ao-49-36-6824>
- [12] Y. Liu, X. Yuan, J. Suo, D. J. Brady, and Q. Dai, "Rank minimization for snapshot compressive imaging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 12, pp. 2990–3006, 2019.
- [13] X. Yuan, "Generalized alternating projection based total variation minimization for compressive sensing," in *2016 IEEE International Conference on Image Processing (ICIP)*, 2016, pp. 2539–2543.

- [14] Y. Fu, Y. Zheng, I. Sato, and Y. Sato, "Exploiting spectral-spatial correlation for coded hyperspectral image restoration," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3727–3736.
- [15] L. Huang, R. Luo, X. Liu, and X. Hao, "Spectral imaging with deep learning," *Light: Science & Applications*, vol. 11, no. 1, p. 61, Mar 2022. [Online]. Available: <https://doi.org/10.1038/s41377-022-00743-6>
- [16] X. Miao, X. Yuan, Y. Pu, and V. Athitsos, "lambda-Net: Reconstruct Hyperspectral Images From a Snapshot Measurement," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 4058–4068.
- [17] Z. Meng, J. Ma, and X. Yuan, "End-to-end low cost compressive spectral imaging with spatial-spectral self-attention," in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII*. Berlin, Heidelberg: Springer-Verlag, 2020, p. 187–204. [Online]. Available: [https://doi.org/10.1007/978-3-030-58592-1\\_12](https://doi.org/10.1007/978-3-030-58592-1_12)
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf)
- [19] Y. Cai, J. Lin, X. Hu, H. Wang, X. Yuan, Y. Zhang, R. Timofte, and L. Van Gool, "Mask-guided spectral-wise transformer for efficient hyperspectral image reconstruction," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun 2022, pp. 17 481–17 490. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR52688.2022.01698>
- [20] —, "Coarse-to-fine sparse transformer for hyperspectral image reconstruction," in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 686–704.
- [21] T. Huang, W. Dong, X. Yuan, J. Wu, and G. Shi, "Deep gaussian scale mixture prior for spectral compressive imaging," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun 2021, pp. 16 211–16 220. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR46437.2021.01595>
- [22] Y. Cai, J. Lin, H. Wang, X. Yuan, H. Ding, Y. Zhang, R. Timofte, and L. V. Gool, "Degradation-aware unfolding half-shuffle transformer for spectral compressive imaging," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 37 749–37 761. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/f621c2ead473ca36763696b712ffda01-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/f621c2ead473ca36763696b712ffda01-Paper-Conference.pdf)
- [23] Y. Dong, D. Gao, T. Qiu, Y. Li, M. Yang, and G. Shi, "Residual degradation learning unfolding framework with mixing priors across spectral and spatial for compressive spectral imaging," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun 2023, pp. 22 262–22 271. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR52729.2023.02132>
- [24] M. Li, Y. Fu, J. Liu, and Y. Zhang, "Pixel adaptive deep unfolding transformer for hyperspectral image reconstruction," in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct 2023, pp. 12 913–12 922. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/ICCV51070.2023.01191>
- [25] L. Song, L. Wang, M. H. Kim, and H. Huang, "High-accuracy image formation model for coded aperture snapshot spectral imaging," *IEEE Transactions on Computational Imaging*, vol. 8, pp. 188–200, 2022.
- [26] L. Wang, Z. Xiong, D. Gao, G. Shi, W. Zeng, and F. Wu, "High-speed hyperspectral video acquisition with a dual-camera architecture," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 4942–4950.
- [27] L. Wang, Z. Xiong, D. Gao, G. Shi, and F. Wu, "Dual-camera design for coded aperture snapshot spectral imaging," *Appl. Opt.*, vol. 54, no. 4, pp. 848–858, Feb. 2015.
- [28] Z. Xiong, L. Wang, H. Li, D. Liu, and F. Wu, "Snapshot hyperspectral light field imaging," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 6873–6881.
- [29] J. Feng, X. Fang, X. Cao, C. Ma, Q. Dai, H. Zhu, and Y. Wang, "Advanced hyperspectral video imaging system using amici prism," *Opt. Express*, vol. 22, no. 16, pp. 19 348–19 356, Aug 2014. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-22-16-19348>
- [30] A. Rouxel, A. Monmayrant, S. Lacroix, H. Camon, and S. Lopez, "Accurate ray-tracing optical model for coded aperture spectral snapshot imagers," *Applied Optics*, vol. 63, no. 7, p. 1828, Feb. 2024.
- [31] Q. Sun, C. Wang, Q. Fu, X. Dun, and W. Heidrich, "End-to-end complex lens design with differentiate ray tracing," *ACM Transactions on Graphics*, vol. 40, no. 4, pp. 1–13, Jul. 2021.
- [32] E. Tseng, A. Mosleh, F. Mannan, K. St-Arnaud, A. Sharma, Y. Peng, A. Braun, D. Nowrouzezahrai, J.-F. Lalonde, and F. Heide, "Differentiable compound optics and processing pipeline optimization for end-to-end camera design," *ACM Trans. Graph.*, vol. 40, no. 2, Jun. 2021. [Online]. Available: <https://doi.org/10.1145/3446791>
- [33] C. Wang, N. Chen, and W. Heidrich, "dO: A Differentiable Engine for Deep Lens Design of Computational Imaging Systems," *IEEE Transactions on Computational Imaging*, vol. 8, pp. 905–916, 2022.
- [34] N. Diaz, J. Ramirez, E. Vera, and H. Arguello, "Adaptive multisensor acquisition via spatial contextual information for compressive spectral image classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 9254–9266, 2021.
- [35] H. Arguello, H. Rueda, Y. Wu, D. W. Prather, and G. R. Arce, "Higher-order computational model for coded aperture spectral imaging," *Appl. Opt.*, vol. 52, no. 10, pp. D12–D21, Apr 2013. [Online]. Available: <https://opg.optica.org/ao/abstract.cfm?URI=ao-52-10-D12>
- [36] I. Choi, D. S. Jeon, G. Nam, D. Gutierrez, and M. H. Kim, "High-quality hyperspectral reconstruction using a spectral prior," *ACM Trans. Graph.*, vol. 36, no. 6, Nov. 2017. [Online]. Available: <https://doi.org/10.1145/3130800.3130810>
- [37] L. Paillet, A. Rouxel, H. Carfantan, A. Monmayrant, and S. Lacroix, "Differentiable chief-ray tracing simulator for coded-aperture spectral imaging," in *Optica Imaging Congress 2024 (3D, AOMS, COSI, ISA, pcAOP)*, ser. COSI, vol. 7. Optica Publishing Group, 2024, p. CTh4B.2.
- [38] A. Rouxel and L. Paillet, "SIMCA," <https://github.com/a-rouxel/simca>, 2023.
- [39] G. Côté, F. Mannan, S. Thibault, J.-F. Lalonde, and F. Heide, "The differentiable lens: Compound lens search over glass surfaces and materials for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023.
- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014.
- [41] R. Zhao, C. Yang, R. T. Smith, and L. Gao, "Coded aperture snapshot spectral imaging fundus camera," *Scientific Reports*, vol. 13, no. 1, p. 12007, Jul 2023. [Online]. Available: <https://doi.org/10.1038/s41598-023-39117-2>
- [42] M. Zhang, L. Wang, and H. Huang, "Sub-pixel dispersion model for coded aperture snapshot spectral imaging," *IEEE Transactions on Computational Imaging*, vol. 9, pp. 1188–1200, 2023.
- [43] J. Park, M. Lee, M. Grossberg, and S. Nayar, "Multispectral Imaging Using Multiplexed Illumination," in *IEEE International Conference on Computer Vision (ICCV)*, Oct 2007.
- [44] L. Wang, C. Sun, Y. Fu, M. H. Kim, and H. Huang, "Hyperspectral image reconstruction using a deep spatial-spectral prior," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8024–8033.
- [45] R. H. Yuhas, A. F. H. Goetz, and J. W. Boardman, "Discrimination among semi-arid landscape endmembers using the spectral angle mapper (sam) algorithm," *Summaries of the Third Annual JPL Airborne Geoscience Workshop*, vol. 1, June 1992.
- [46] E. H. Linfoot, "Information theory and optical images," *J. Opt. Soc. Am.*, vol. 45, no. 10, pp. 808–819, Oct 1955. [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=josa-45-10-808>
- [47] J. O'Sullivan, R. Blahut, and D. Snyder, "Information-theoretic image formation," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2094–2123, 1998.
- [48] A. Ashok and M. A. Neifeld, "Information-based analysis of simple incoherent imaging systems," *Opt. Express*, vol. 11, no. 18, pp. 2153–2162, Sep 2003. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-11-18-2153>