



OpenSpyrit: an ecosystem for open single-pixel hyperspectral imaging

GUILHERME BENETI MARTINS,¹ LAURENT MAHIEU-WILLIAME,¹
THOMAS BAUDIER,¹ AND NICOLAS DUCROS^{1,2,*} 

¹Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, LYON, France

²Institut universitaire de France (IUF), France

*nicolas.ducros@creatis.insa-lyon.fr

Abstract: This paper describes OpenSpyrit, an open access and open source ecosystem for reproducible research in hyperspectral single-pixel imaging, composed of SPAS (a Python single-pixel acquisition software), SPYRIT (a Python single-pixel reconstruction toolkit) and SPIHIM (a single-pixel hyperspectral image collection). The proposed OpenSpyrit ecosystem responds to the need for reproducibility and benchmarking in single-pixel imaging by providing open data and open software. The SPIHIM collection, which is the first open-access FAIR dataset for hyperspectral single-pixel imaging, currently includes 140 raw measurements acquired using SPAS and the corresponding hypercubes reconstructed using SPYRIT. The hypercubes are reconstructed by both inverse Hadamard transformation of the raw data and using the denoised completion network (DC-Net), a data-driven reconstruction algorithm. The hypercubes obtained by inverse Hadamard transformation have a native size of $64 \times 64 \times 2048$ for a spectral resolution of 2.3 nm and a spatial resolution that is comprised between 182.4 μm and 15.2 μm depending on the digital zoom. The hypercubes obtained using the DC-Net are reconstructed at an increased resolution of $128 \times 128 \times 2048$. The OpenSpyrit ecosystem should constitute a reference to support benchmarking for future developments in single-pixel imaging.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Spectral imaging is a major tool of modern science, with applications in astronomy, environmental monitoring, food processing, agriculture, and biomedical imaging [1,2]. The pushbroom and filter-based methods are scanning techniques that require multiple measurements to acquire a full (x, y, λ) hypercube. Pushbroom methods acquire one (y, λ) slice at a time and require scanning along the x -axis [3]. Filter-based setups acquire an (x, y) image for one spectral band, with the full hypercube obtained from a sequence of measurements by rotating a filter wheel or monitoring electronically tunable filters [4]. Both pushbroom and filter-based approaches suffer from low optical throughput as only a small part of the hypercube is measured at a time. Moreover, the spatial (i.e., for pushbroom) or spectral (i.e., for filters) resolutions are linear in proportion to the number of measurements, and hence they are either slow or low resolution. These limitations have led to computational snapshot methods that rely on algorithms that reconstruct a hypercube from a few raw measurements [5]. Coded aperture snapshot spectral imagers and their different variants exploit a diffractive element with a programmable mask, such that each raw measurement gives access to an oblique projection of the hypercube [6]. Miniature ultra-spectral imaging uses a liquid crystal phase retarder to multiplex the spectral domain [7]. The spectral DiffuserCam is a multispectral filter array where the optics are replaced by a diffuser, such that each pixel on the sensor can ‘see’ the whole field of view [8]. However, snapshot imagers suffer from an inherent trade-off between the spatial and the spectral dimensions.

Hyperspectral single-pixel imaging is a generalization of single-pixel imaging whereby a reconstruction algorithm is used to recover high-spectral-resolution hypercubes from a set of

spectra obtained using a set of spatial light patterns [9–14]. Compared to snapshot imagers, this approach requires more measurements; however, it leads to an excellent spectral resolution, as the spectral dimension is not multiplexed. Hyperspectral single-pixel imaging traces back to the concept of Hadamard spectroscopy [15] that introduced Hadamard-coded measurements to obtain a signal-to-noise ratio boost known as Fellgett's effect [16]. This approach has received renewed interest thanks to the advent of compressed sensing theory that allows the number of measurements to be drastically reduced [17]. Interest has been further increased with the advent of deep learning that enables fast reconstruction while outperforming handcrafted prior-based methods [18]. The number of algorithms that exploit deep learning for single-pixel reconstruction has continuously increased (see [19–25] to cite only a few relevant works). There is now a wide variety of single-pixel systems and methods. Different systems are obtained by integrating different light modulators, detectors, and synchronization strategies, while different single-pixel methods are obtained by considering different types of light patterns, data preprocessing steps, and reconstruction algorithms. However, there is no common dataset and platform for comparing and testing the various methods. Concerns exist about the characterization of the performance (e.g., resolution, sensitivity, robustness to noise) of these computational systems compared to traditional ones (e.g., pushbroom). This includes concerns about data-driven algorithms whose results are often difficult to reproduce and compare due to the hyperparameters and stochastic nature of the training phase.

Here, we propose the OpenSpyrit ecosystem which addresses the need for reproducibility and benchmarking in single-pixel imaging by providing open data and open software. Our new single-pixel datasets include both the raw measurements and the corresponding hypercubes reconstructed by two algorithms. Our acquisition and reconstruction software not only implement a particular methodology but also provide tools that can be tailored to different user-specific configurations. We have acquired a collection of hypercubes over 2,048 spectral channels using a single-pixel imaging device whose spectral and spatial resolutions were characterized. Our imaging device relies on a digital micromirror device (DMD) that displays a sequence of Hadamard patterns. It acquires a spectrum for each pattern using a commercial spectrometer. While the hypercubes can be recovered directly by inverse Hadamard transformation of the raw spectra, we also consider an explainable reconstruction method based on deep learning [21]. Our datasets - SPIHIM [26] - are made publicly available following FAIR (findability, accessibility, interoperability, reusability) principles [27]. To the best of our knowledge, SPIHIM is the first FAIR open-access data collection for hyperspectral single-pixel imaging. We share the raw data together with the Hadamard inverse transformed hypercubes and the hypercubes reconstructed thanks to deep learning. Our acquisition software is released via the SPAS package [28]. Our reconstruction method is implemented in the Python package SPYRIT [29], which allows the method to be retrained from scratch or to evaluate the corresponding Pytorch models that we make available. The goal of OpenSpyrit is not to compare all existing single-pixel configurations and methodologies, which would not be possible, but to provide novel tools that enable fair benchmarking and reproducible research for all future developments in single-pixel imaging.

2. Methods

The computational framework implemented in the OpenSpyrit ecosystem is depicted in Fig. 1. To acquire a 3D hypercube using a 2D sensor, we acquire multiple pixels at the same time by shaping the light with a DMD. After a sequence of spectra has been taken by using different DMD patterns, we feed the raw spectra into a deep reconstruction algorithm that recovers the hypercube.

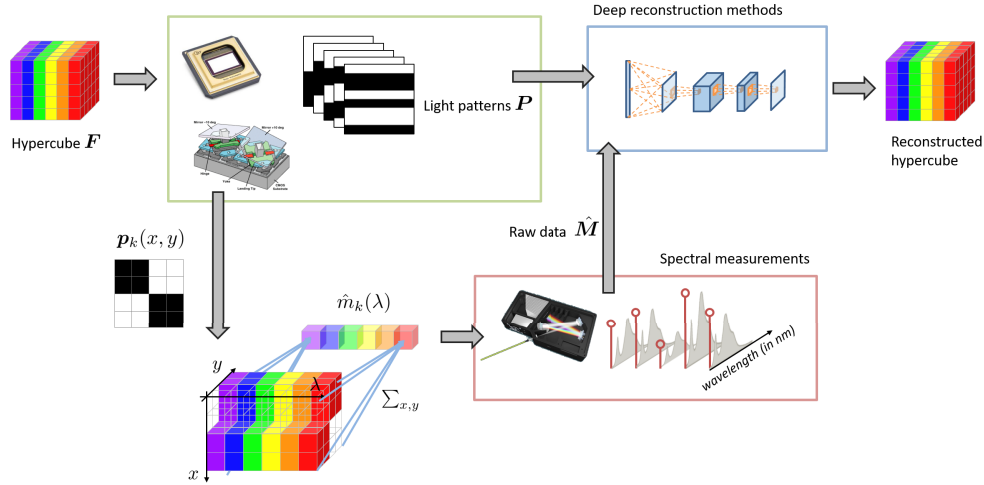


Fig. 1. Hyperspectral single-pixel imaging principle. The hypercube $F \in \mathbb{R}^{N \times \Lambda}$ is sent to a compact spectrophotometer via a digital micromirror device (DMD). A sequence of $2K$ light patterns $P \in \mathbb{R}^{2K \times N}$ is uploaded onto the DMD, leading to the measurement of the $2K$ raw spectra $\hat{M} \in \mathbb{R}^{2K \times \Lambda}$. A reconstruction method is then used to reconstruct the hypercube F from the raw spectra \hat{M} .

2.1. Image formation model

Let $\hat{M} \in \mathbb{R}^{2K \times \Lambda}$ represent the raw measurements, where $2K$ is the number of DMD patterns and Λ the number of spectral channels provided by the spectrometer. Let $P \in \mathbb{R}^{2K \times N}$ be the matrix that contains the DMD patterns, where N is the number of (spatial) pixels in each pattern and $F \in \mathbb{R}^{N \times \Lambda}$ represents the 3D hypercube. We model the acquisition process as linear measurements corrupted by Poissonian-Gaussian noise [30]

$$\hat{M} \sim g\mathcal{P}(PF) + \mathcal{N}(\mu_{\text{dark}}, \sigma_{\text{dark}}^2), \quad (1)$$

where \mathcal{P} and \mathcal{N} are the Poisson and Gaussian distributions, g represents the system gain (in counts/electron), μ_{dark} is the dark current (in counts), and σ_{dark} is the dark noise (in counts).

We choose P as Hadamard patterns, which maximizes the signal-to-noise ratio of the reconstructed hypercube with respect to the additive Gaussian noise [15]. To handle the negative values in P , the light patterns are split into positive and negative parts to be uploaded onto the DMD [31]. In notations, we have

$$P = \begin{bmatrix} P_+ \\ P_- \end{bmatrix}, \quad (2)$$

where $P_+ \in \mathbb{R}_+^{K \times N}$ and $P_- \in \mathbb{R}_+^{K \times N}$ are the positive and negative parts of Hadamard patterns respectively, i.e., $P_+ - P_- = SH$, where $H \in \mathbb{R}^{N \times N}$ is the Walsh-Hadamard basis and $S \in \{0, 1\}^{K \times N}$ is a subsampling matrix that retains some of the rows of H . In the following, we denote the retained Hadamard patterns by $H_{\downarrow} = SH$. Different subsampling strategies (e.g., random, low frequency, high variance) have been investigated (see [32] for an overview).

We finally preprocess the raw measurements $\hat{M} = \begin{bmatrix} \hat{M}_+ \\ \hat{M}_- \end{bmatrix}$ to compensate for splitting

$$M = \hat{M}_+ - \hat{M}_-, \quad (3)$$

where $\hat{M}_+ \in \mathbb{R}^{K \times \Lambda}$ and $\hat{M}_- \in \mathbb{R}^{K \times \Lambda}$ correspond to the measurements obtained with the positive and negative patterns, respectively. Therefore, the preprocessed measurements M are Hadamard

coefficients, in the sense that $\mathbb{E}(\mathbf{M}) = \mathbf{H}_\downarrow \mathbf{F}$, where \mathbb{E} denotes the expectation. Note that the problem is separable across the spectral dimension, i.e., $\mathbb{E}(\mathbf{m}_\lambda) = \mathbf{H}_\downarrow \mathbf{f}_\lambda$, $1 \leq \lambda \leq \Lambda$, where $\mathbf{m}_\lambda \in \mathbb{R}^K$ and $\mathbf{f}_\lambda \in \mathbb{R}^N$ are the λ -th column of \mathbf{M} and \mathbf{P} , respectively. Therefore, the spectral resolution of the hypercube is given directly by the spectral resolution of the spectrometer, while its spatial resolution depends only on the light patterns and our ability to recover \mathbf{f}_λ from \mathbf{m}_λ .

2.2. Image reconstruction

The hypercube can be reconstructed in the least squares sense as

$$\mathbf{F} = \frac{1}{N} \mathbf{H}_\downarrow^\top \mathbf{M}. \quad (4)$$

In the case $K = N$, the pseudo inverse $\frac{1}{N} \mathbf{H}_\downarrow^\top$ is the inverse of \mathbf{H} .

We also propose to reconstruct each λ -slice of the hypercube independently by using a data-driven algorithm. We consider the denoised completion network (DC-Net) [33] that computes the two steps

$$\bar{\mathbf{f}}_\lambda = \mathcal{G}_{\text{dc}}(\mathbf{m}_\lambda), \quad \text{and} \quad (5a)$$

$$\mathbf{f}_\lambda = \mathcal{D}_\theta(\bar{\mathbf{f}}_\lambda), \quad (5b)$$

where \mathcal{G}_{dc} represents the denoised completion step and \mathcal{D}_θ represents a neural network with parameters θ . The denoised completion network is given by $\mathcal{G}_\theta = \mathcal{D}_\theta \circ \mathcal{G}_{\text{dc}}$.

The operator \mathcal{G}_{dc} is chosen as the linear estimator achieving minimum mean squared error under Gaussian assumptions. It is given by

$$\mathcal{G}_{\text{dc}}(\mathbf{m}) = \frac{1}{N} \mathbf{H}_\downarrow^\top \begin{bmatrix} \mathbf{I}_M \\ \boldsymbol{\Sigma}_{21} \boldsymbol{\Sigma}_1^{-1} \end{bmatrix} \boldsymbol{\Sigma}_1 (\boldsymbol{\Sigma} + \boldsymbol{\Sigma}_1)^{-1} \mathbf{m}, \quad (6)$$

where $\mathbf{I}_M \in \mathbb{R}^{M \times M}$ is the identity matrix, $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_{21}$ are blocks of the covariance matrix of $\mathbf{H}\mathbf{f}$ and $\boldsymbol{\Sigma}$ is the noise covariance, which can be estimated as detailed in [21]. While the denoised completion operator \mathcal{G}_{dc} is kept fixed during training, we optimize the parameters of the convolutional neural network in a supervised manner

$$\arg \min_{\theta} \sum_{\ell} \|\mathbf{f}^{(\ell)} - \mathcal{G}_\theta(\mathbf{m}^{(\ell)})\|^2, \quad (7)$$

where $\{\mathbf{f}^{(\ell)} \in \mathbb{R}^N\}_{1 \leq \ell \leq L}$ is an image database and $\{\mathbf{m}^{(\ell)} \in \mathbb{R}^K\}_{1 \leq \ell \leq L}$ are the associated measurements computed according to Eq. (1) and Eq. (3).

2.3. SPAS: single-pixel acquisition software

SPAS [28] is an open source python package for single-pixel acquisition, which has been tested to control a DMD 4100 (0.7" XGA VIS, ViALUX) and a spectrometer (Avantes AvaSpec-ULS2048CL-EVO). It therefore requires the ALP4lib [34] package for DMD control and the MSL-Equipment [35] package for spectrometer control. SPAS provides simple functions for the initialization of the DMD and spectrometer, for the upload of a sequence of patterns into the DMD memory, and for the acquisition that is triggered by an external signal provided by the DMD. Each pattern is displayed on the DMD during a given illumination time and the external trigger is sent to the spectrometer for synchronization. During acquisition, the spectra are stored in the spectrometer's internal memory and are transferred to the computer via a callback function, which allows other tasks such as image reconstruction to be run in parallel.

SPAS also provides several functions for the visualization of the hypercubes (e.g., spectral slicing, spectral binning, or RGB representation based on CIE color matching functions [36]), in addition to the implementation of a simple interface to the more advanced reconstruction functions of the SPYRIT package (see Section 2.5).

2.4. SPIHIM: single-pixel hyperspectral imaging data collection

SPIHIM [26] is a FAIR [27], open-access collection of hypercubes acquired using the principle of single-pixel imaging. The collection is stored in the PILoT warehouse managed by Girder, an open source web-based data management platform. Each acquisition is associated with a uniquely identified folder which contains the raw measurements (`*_spectraldata.npz`), the hypercube reconstructed by inverse Hadamard transformation (`*_had_reco.npz`), the hypercube reconstructed using our neural network (`*_nn_reco.npz`), some metadata (`*_metadata.json`) and an overview folder with several spatial and spectral representations of the hypercube. The metadata provide information on the experimental conditions as well as on the acquisition parameters. The metadata together with the raw data and hypercubes reconstructed by inverse Hadamard transformation should greatly facilitate the implementation of benchmarking and reproducibility. The collection currently contains 140 hypercubes acquired from 15 different objects: "Star Sector", "Cat", "PpIX", "Color checker", "USAF", "Colored Star Sector", "No object", "Tomato slice", "Horse", "White spot", "Thorlabs box", "Blob", "Tree leaf", "Lamb brain", "Apple". For further details about the objects, please see Section 1A of [Supplement 1](#). We also provide an exhaustive list and detailed description of the acquisitions made from these objects (see Tables S1-S15 in [Supplement 1](#)).

2.5. SPYRIT: single-pixel reconstruction toolkit

SPYRIT is a Python package for single-pixel image reconstruction [29] that relies on PyTorch [37] and can be easily installed using the pip package management system. To handle different problems in a simple and effective manner, SPYRIT is divided into four core submodules that implement the measurement operator, noise model, preprocessing step, and reconstruction algorithm via classes that inherit from `nn.module` (see Fig. 2). The pipeline from an image acquisition to its reconstruction can be seen as a neural network, for which all PyTorch functionalities (e.g., data loaders, network architectures, loss functions, optimizers, etc) can be exploited. SPYRIT already includes the specific measurement operator, noise model, preprocessing step, and reconstruction step corresponding to Eq. (1), (3) and (5), as well as base classes that SPYRIT users can tailor to their particular problem. The source code of the SPYRIT package is available on GitHub and benefits from continuous integration with automated release of Python packages, dedicated benchmarks for regression testing, and integrated documentation. The software is licensed under LGPL-3.0: it can be used and modified by anyone for private, public or commercial use. The SPYRIT package is complemented by a companion Github repository [38] that contains script examples that rely on SPYRIT. In particular, the scripts that generate the figures displayed in this manuscript and [Supplement 1](#) are available in `/2022_OE_spyrit2/`.




Fig. 2. Overview of the four SPYRIT core modules combined in a typical training pipeline: `meas` represents the measurement operator, `noise` the noise model, `prep` the preprocessing step, and `recon` the reconstruction algorithm. The modularity of SPYRIT allows a variety of computational problems to be handled in an effective and simple manner.

3. Experiments

3.1. Experimental setup

Our setup is composed of an illumination arm, a DMD, and a light collection arm as depicted in Fig. 3. The illumination arm is composed of a white LED lamp (Thorlabs LIUCWHA) and a bi-telecentric lens system (Edmund Optics TECHSPEC Large Format Telecentric 62902, magnification 0.9x) that forms the image of the object in the active plane of a DMD (ViALUX GmbH DLP V-700, 1024 x 768 micromirrors, 13.7 μm pitch). The DMD is made of a matrix of microscopic mirrors that can be individually tilted to either $+24^\circ$ (ON state) or -24° (OFF state) according to spatial light patterns. The light collection arm, placed at $+24^\circ$ with respect to the illumination arm, holds a 35 mm focal length bi-convex lens, a MAP104040-B Matched Achromatic Lens Pair (both focal lengths are 40 mm), and an objective lens (x20, NA = 0.35) that focuses light at the entrance of an optical fiber (1500 μm core diameter, NA = 0.39, FT1500 UMT) connected to a compact spectrometer (Avantes AvaSpec-ULS2048CL-EVO, $\Lambda = 2048$ spectral channels, 515–750 nm, entrance slit of 200 μm , 1200 lines/mm grating). The setup, which is supported by a cage system, is lightweight and transportable. The setup depicted in Fig. 3 is referred to as version v1.1 in the SPIHIM collection. The variants v1.2, v1.3 and v1.3.1 differ with respect to the optical elements in the illumination and collection arms of the setup. Additional details are provided in the SPIHIM collection [26].




Fig. 3. Acquisition system. Light source (L), sample (S), telecentric lens (TL), digital micromirror device (DMD), bi-convex lens (CL), achromatic lens pair (LP), objective lens (OL), optical fiber (OF), spectrometer (SP), and instrumentation computer (PC). The green arrows indicate the communication workflow between the computer, the DMD, and the spectrometer. The blue arrows indicate the light path.

The integration time of the spectrometer is chosen as equal to the illumination time. A dead time of 44 μs , referred to as dark phase, is necessary for the DMD to tilt its micromirrors according to the next pattern. Another dead time of 356 μs is necessary for the spectrometer to flush its buffer and prepare a new acquisition. This leads to an acquisition time per pattern equal to the spectrometer integration time plus the longest dead time. The total time for the acquisition of a hypercube is $T = 2K(\Delta t + \delta t)$, where Δt represents the integration time and δt represents the dead time. While the integration time can be chosen by the user, the dead time is imposed by the device. Note that the dead time of the spectrometer is much longer than the smallest illumination time allowed by the DMD which cannot be operated at its maximum frequency (22 kHz corresponding to 45 μs). Our acquisitions are typically made with an integration time

of 1 ms. Therefore, the fully sampled acquisition of an image of $N = 64 \times 64$ pixels requires $2K = 2N = 8,192$ patterns $\times 1.4$ ms ≈ 11.5 s.

3.2. Experimental data for setup characterization

As a reference image, we consider the image of a cat taken from the STL-10 [39] test set. The cat image is printed on a plastic sheet on which we superimpose a linear variable filter (Ocean Optics, LVF-HL, see Fig. 4). For the purpose of reproducibility, we characterize the spatial and spectral resolution of the setup. To do so, we consider five objects from the SPIHIM collection. To determine the spatial resolution, we consider the Star Sector resolution target (Thorlabs, R1L1S2P, see top row of Fig. 5 and Fig. 6) and the USAF resolution target (Edmund, USAF 1951 38256, see bottom row of 5 and Fig 6). Both samples are illuminated in transmission using a cold white LED array light source (Thorlabs, LIUCWHA). To determine the spectral resolution, we consider the light spots of a Mercury-Argon calibration lamp (Ocean Optics HG-1 with characteristic peaks at 546, 577, 579, 697, 707, 727, and 738 nm, see Fig. 7). Finally, we consider an image of a tomato slice (see Fig. 8) to evaluate the capacity of the neural network to reconstruct the hypercube as a function of the acceleration rate. A tomato slice has smooth optical contrast compared to calibration targets that have high symmetry and sharp edges where reconstruction can fail for an undersampled acquisition. Details about these acquisitions can be found in Supplement 1 (see Tables S1, S2, S5, S7 and S8).




Fig. 4. STL-10 cat hypercube acquisition with a linear variable filter. The full hypercube is binned spectrally for display (seven bins in the range 544–670 nm, bin widths ~ 19 nm, central wavelengths: 553.7, 572.6, 591.1, 609.3, 627.1, 644.5, and 661.5 nm). The colorbars show intensities in counts/pixel. The image on the bottom right is an RGB representation of the full hypercube. Acquisition: $K = N = 4,096$ patterns, $\times 1$ zoom, integration time of 1 ms/pattern; reconstruction by means of Eq. (4) with $N = 4,096$.

Irrespective of the imaging configuration, all the patterns of a 64×64 Hadamard basis are acquired, resulting in a total of $K = N = 4,096$ Hadamard patterns split into 8,192 positive and negative patterns. We also consider accelerated acquisitions for which $K < 4,096$ patterns are acquired. The fully sampled datasets can be downsampled *a posteriori* to simulate an accelerated acquisition with different acceleration factors. Typically, the 64×64 Hadamard patterns are resized to fill the largest square region possible on the DMD, which corresponds to 768×768 micromirrors. In this case, each pixel of the Hadamard patterns corresponds to an area of 12×12 micromirrors. However, it is also possible to display the patterns on smaller fields of view, acting as a hardware zoom that is independent of the optical components of the acquisition setup. We




Fig. 5. Resolution targets acquired with different zooms. Top row: Star Sector; bottom row: USAF. Zoom increases from left to right: $\times 1$, $\times 3$, $\times 6$, and $\times 12$; $K = N = 4,096$ patterns; the integration time increases with the zoom: 1, 9, 36, and 144 ms/pattern, from left to right. All hypercubes are reconstructed using Eq. (4). The displayed images are obtained by summing the hypercubes in the 550-590 nm range.

consider six DMD-based hardware zooms: $\times 1$, $\times 2$, $\times 3$, $\times 4$, $\times 6$, and $\times 12$, which correspond to patterns with a pixel size of 12, 6, 4, 3, 2, and 1 micromirrors, respectively. The higher the zoom factor, the lower the photon counts. To obtain measurements with similar signal-to-noise ratios, we chose the integration time depending on the zoom, as indicated in Table 1. For the $\times 1$ zoom, we set the integration time to 1 ms/pattern for the STL-10 cat and 17.4 ms/pattern for the Mercury-Argon calibration lamp. For the Star Sector and USAF targets that we image at zooms $\times 1$, $\times 2$, $\times 3$, $\times 4$, $\times 6$, and $\times 12$, we choose integration times of 1, 4, 9, 16, 36, and 144 ms/pattern, respectively. For the tomato slice, we choose 4 ms/pattern for $\times 2$ zoom and 144 ms/pattern for $\times 12$ zoom.

Table 1. Integration times for the different zooms and samples. Times are given in ms/pattern; ‘n.a.’ indicates that a dataset is not available. The corresponding images are available in Supplement 1 (see Tables S1, S2, S5, S7 and S8). Click on the integration times to access the raw data via a unique identifier.

Zoom	$\times 1$	$\times 2$	$\times 3$	$\times 4$	$\times 6$	$\times 12$
Pixel size (μm)	182.4	91.2	60.8	45.6	30.4	15.2
STL-10 cat	1	n.a.	n.a.	n.a.	n.a.	n.a.
Star Sector	1	4	9	16	36	144
USAF	1	4	9	16	36	144
Mercury-Argon lamp	17.4	n.a.	n.a.	n.a.	n.a.	n.a.
Tomato slice	1	4	n.a.	n.a.	36	144

3.3. Training of the DC-Net

We train a DC-Net using SPYRIT for a number of measurements K of 4096, 2048, 1024, and 512 and images of size $N = 128 \times 128$, contrary to the original implementation [21] that considers smaller images of size $N = 64 \times 64$. To do so, we consider the $L = 100,000$ images that correspond to the ‘test’ subset of the ImageNet dataset [40]. The original images are randomly

cropped to 128×128 and are normalized between -1 and 1 . The image domain denoiser \mathcal{D}_θ is a U-Net with three downsampling steps separated by a max pooling operation and three upsampling steps separated by transposed convolutions. The number of filters in the contracting path is increased from 16 to 32, then to 64. Each layer is separated by a ReLU and batch normalization layer. Our U-Net has a total of 499,985 learnable parameters. Here again, this differs from the original implementation [21] that considers a simple three-layer convolutional neural network with 28,993 learnable parameters. We solve Eq. (7) using the ADAM optimizer [41], with an initial learning rate of 10^{-3} , which is halved every 10 epochs, for a maximum of 30 epochs. The training phase took about 300 minutes on a NVIDIA GeForce RTX 2080 Ti graphic card in the case of $K = 512$. The DC-Nets trained for the four number of measurements are available in the `/soft/model/` folder of SPIHIM.

4. Results

4.1. Acquisition example

To show the spatial and spectral capabilities of our hyperspectral camera, we first consider the STL-10 cat image with a linear variable filter in front of it. The integration time was set at 1 ms/pattern leading to a total acquisition time of 11.5 s. Figure 4 shows the obtained hypercube to which we apply spectral binning to facilitate its display. We compute seven bins within the 544–670 nm range with a bin width of ~ 19 nm (central wavelengths: 553.7, 572.6, 591.1, 609.3, 627.1, 644.5, and 661.5 nm). We also provide an RGB representation as recommended by the CIE [36]. Details about this acquisition can be found in [Supplement 1](#) (see Table S2).

Each bin displays a different bandpass window that is selected by the linear filter. As expected, the bandpass window translates diagonally within the field of view, from the top left corner to the bottom right corner, as the central wavelength increases. This is also visible on the RGB representation that displays the color palette starting with the green color in the top left corner to the red color in the bottom right corner. The STL-10 cat is visible in the background.

4.2. Spatial resolution and DMD-based zoom

We evaluate the spatial resolution of our system by imaging two calibrated resolution targets: the Star Sector and the USAF target. The Star Sector is composed of 36 black bars distributed around 360° . The USAF target is composed of bar groups with decreasing bar spacing and length. In Fig. 5, we display the images obtained for both targets at four different zooms ($\times 1$, $\times 3$, $\times 6$, and $\times 12$), after summation in the 550–590 nm range. The integration time was set at 1, 9, 36, and 144 ms/pattern respectively leading to a total acquisition time of 11.5, 77.0, 298.2, and 1,183.0 s respectively. Details about these acquisitions can be found in [Supplement 1](#) (see Tables S1 and S5).

For both targets, we first establish the spatial resolution in pixels. Then, we convert it to line pairs per millimeter (lp/mm) to account for the optical magnification. For the Star Sector, we determine the system resolution as the smallest radius of a circular profile for which consecutive black bars appear to touch. For the USAF, we determine the system resolution as the smallest distinguishable bar group vertically and horizontally [42]. We also report the theoretical spatial resolution computed as $1/(2\Delta x)$, where Δx is the image pixel size in millimeters. The image pixel size depends linearly on the zoom, given the DMD pixel size and the telecentric lens magnification. We obtain 182.4, 91.2, 60.8, 45.6, 30.4, and $15.2 \mu\text{m}$ for $\times 1$, $\times 2$, $\times 3$, $\times 4$, $\times 6$, and $\times 12$ zooms, respectively (see Table 1).

In Fig. 6, we plot the spatial resolution as a function of the zoom, considering six different zooms that correspond to six independent acquisitions. We observe that spatial resolutions obtained from both the USAF and Star Sector targets are in good agreement with theoretical

values computed from the pixel size only. This indicates that our system is limited only by the pixel size and that the DMD-based hardware zoom is not associated with undesirable blur.

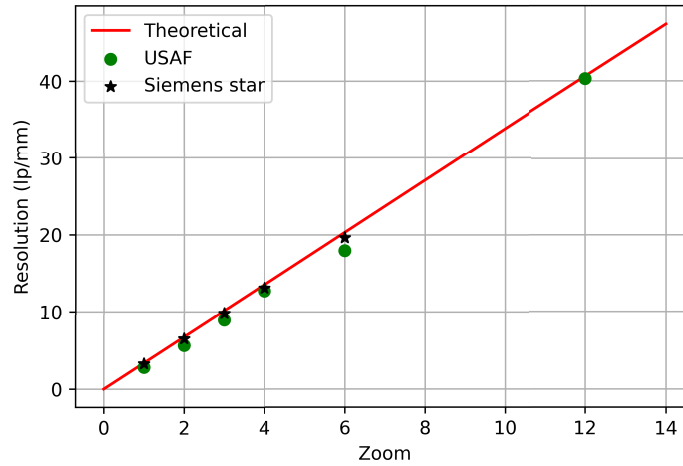


Fig. 6. Spatial resolution as a function of the zoom. The red line is the theoretical resolution calculated from the pixel size; the green dots represent the resolution measured from the USAF target; the black stars represent the resolution measured from the Star Sector target. The spatial resolution is given in line pairs/millimeters (lp/mm) and evaluated from the images displayed in Fig. 5.

4.3. Spectral resolution at different locations

We evaluate the spectral resolution by imaging a Mercury-Argon calibration lamp positioned in the object plane. We place the lamp at three different positions to create light spots at different locations in the field of view and acquire a hypercube for each spot. Figure 7 shows the superposition of the three lamp spots ($2K = 8,192$ patterns per acquisition, integration time $\Delta t = 17.4$ ms/pattern). Details about this acquisition can be found in [Supplement 1](#) (see Table S7).

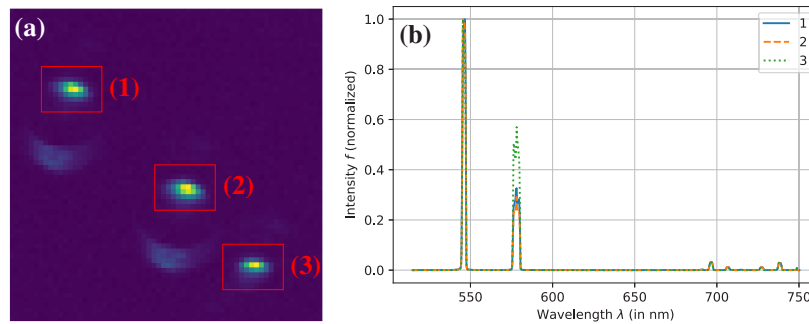


Fig. 7. Spectral resolution at different spatial locations in the field of view. (a) Image of a three-spot of the Mercury-Argon calibration lamp. (b) Spectra of each of the light spots indicated in (a). Integration time: 17.4 ms/pattern. The image in (a) is obtained by summing the hypercube along the spectral dimension. The spectra in (b) are obtained by summing all pixels within the red rectangles displayed in (a). These results confirm that the spectral resolution of our device is directly given by the spectral resolution of the spectrometer and that the spectral response of our system is spatially invariant.

For all three acquisitions we observe a central bright spot corresponding to the position of the light source (Fig. 7(a)). For each spot, we sum the contributions of all pixels in the respective red rectangles, obtaining the spectra indicated in Fig. 7(b). For the three acquisitions, we recover a spectrum that consists of the emission lines of mercury ($\lambda < 650$ nm) and argon ($\lambda > 650$ nm). In the following, we consider the peaks at 546, 697, 707, 727, and 738 nm. We measure the full width at half maximum of all peaks, for all spot locations, and obtain spectral resolutions between 2.15 nm and 2.30 nm. These spectral resolutions are in excellent agreement with the theoretical spectral resolution of the spectrometer that is 2.3 nm, confirming that the spectral resolution of our device is directly given by the spectral resolution of the spectrometer. We observe no spectral degradation that originates from components before the spectrometer (e.g., DMD or focusing optics). Note that the peak at 578 nm results from the observation of the mercury emission doublet at 577 and 579 nm, which cannot be resolved. We also find that the amplitude of the different peaks, except the doublet, are the same for the three spot locations, which indicates that the spectral response of our system is spatially invariant.

4.4. Reconstructions for different values of K and N

We evaluate our ability to reconstruct images from different amounts of measurements considering three samples: the Star Sector target, the USAF target, and a tomato slice. We consider the Star Sector and USAF targets imaged using $\times 12$ zoom, whereas for the tomato slice, both the $\times 2$ and $\times 12$ zooms were considered. Details about these acquisitions can be found in [Supplement 1](#) (see Tables S1, S5 and S8). For each sample, we consider $K = 4,096$ measurements, $K = 1,024$ measurements, and $K = 512$ measurements. The case $K = 4,096$ measurements corresponds to a full basis scan with 64×64 Hadamard patterns. The cases $K = 1,024$ and $K = 512$ measurements are obtained by subsampling the full basis scan acquisition. We retain the patterns that lead to the measurements with the highest energy (i.e., we chose \mathbf{S} such that $\mathbb{E}(\|\mathbf{SHf}\|_2^2)$ is maximal), as suggested in [43]. We reconstruct the hypercubes in the least square sense using Eq. (4) (see first column of Fig. 8) as well as with the data driven DC-Net given by Eq. (5) (see second, third, and fourth columns of Fig. 8). All reconstructed images are obtained by keeping a single spectral channel at $\lambda = 579$ nm. The Hadamard matrices are constructed such that the set of patterns at a given resolution is included in the set of patterns at higher resolutions. For instance, a full basis scan at resolution 64 (i.e., $K = 4,096$ measurements taken from patterns of size $N = 64 \times 64$) is equivalent to an acquisition at resolution 128 (i.e., $K = 16,384$ measurements taken from patterns of size $N = 128 \times 128$) undersampled by a factor of four. To illustrate this, we reconstruct the USAF target at resolution 64 (see second row of Fig. 8) and at resolution 128 (see third row of Fig. 8). More hypercubes reconstructed at resolution 128 can be found in Section 2 of [Supplement 1](#) where we consider fifteen different objects and four different spectral channels.

The least-squares reconstruction is independent of the choice for N and depends on K only, which can be seen by comparing the first image of the second row and the first image of the third row. On the other hand, the reconstruction quality of the data-driven algorithm depends on the size of the images considered during the training phase. Training the DC-Net with 128×128 images rather than with 64×64 images improves the quality of the reconstruction significantly. This can be observed by comparing the images of the second row of Fig. 8 to those of the third row of Fig. 8. In particular, the bars of the elements 5 and 6 of the group 4 can be resolved in the deep reconstruction, but not in the least-squares reconstruction (see the red rectangles in Fig. 8). This improvement is observed even if no such piecewise-constant binary images are present in the ImageNet database used for training. As expected, lowering the number of measurements leads to a loss of spatial resolution, which can be evaluated from the reconstructions of the resolution targets (see second, third, and fourth rows of Fig. 8). The lower the number of measurements, the higher the loss. For the Star Sector, the degradation of the spatial resolution appears as a blurred region in the center of the target, where high spatial frequency structures are present. We also




Fig. 8. Image reconstructions for different samples, number of measurements, and reconstruction methods. First row: subsampling masks for $N = 128 \times 128$; second and third rows: USAF $\times 12$ zoom; fourth row: Star Sector $\times 12$ zoom; fifth row: tomato slice $\times 2$ zoom; sixth row: tomato slice $\times 12$ zoom. First column: $K = 4,096$ measurements, least-squares reconstruction; second column: $K = 4,096$ measurements, denoised completion network (DC-Net) reconstruction; third column: $K = 1,024$ measurements, DC-Net reconstruction; fourth column: $K = 512$ measurements, DC-Net reconstruction. The least-squares reconstructions are obtained using Eq. (4), while the DC-Net reconstructions are given by Eq. (5). We have set $N = 128 \times 128$ in all rows, except in the second where $N = 64 \times 64$. All images correspond to the spectral channel $\lambda = 579$ nm.

observe this effect in the tomato slice images (see fifth and sixth row of Fig. 8). However, as fewer high frequencies are present, the degradation appears more limited.

We also reconstruct images from fifteen different samples for $K = 4096$, 1024 and 512 raw measurements using our DC-Net (See Section 2 of [Supplement 1](#)). For each sample and measurement number, we consider eight images corresponding to four different spectral *channels* and four different spectral *bins* that are obtained by summation of the spectral channels. The resulting 360 reconstructions available constitute a solid reference for future benchmarking.

5. Discussion

Images collected by hardware-driven commercial systems are relatively abundant. Among the main hyperspectral databases, the AVIRIS collection (e.g., see the NASA website [44]) contains thousands of hypercubes over 214 spectral channels taken from airborne platforms. The BGU ICVL collection [45] currently contains 201 hypercubes over 519 spectral channels corresponding to indoor and outdoor scenes taken from the ground. Very few databases offer hypercubes with more than a thousand spectral bands. In contrast to existing datasets, our computational device provides access to 2,048 spectral channels. Our imaging system natively acquires a $64 \times 64 \times 2048$ hypercube with a spectral resolution of 2.3 nm, while the spatial resolution can be adjusted between 182 μm and 15 μm using a DMD-based hardware zoom that can achieve a $\times 12$ magnification with no modification of the optical components. The setup characterization demonstrates that the spectral resolution is not affected by the micromirror matrix of the DMD and is equal to that of the commercial spectrometer, while the spatial resolution is approximately driven by the minimal group of micromirrors. As for the optical zoom, the higher the magnification, the lower the photon flux. To account for this effect, the images at higher zooms have been acquired for a longer duration (see Fig. 5), with a scaling factor equivalent to the zoom squared. Our system maintains a high spectral resolution for a lower price than currently available hyperspectral cameras with the same spectral resolution.

The shortest integration time that we consider is 1 ms/pattern, leading to a total acquisition time of 11.5 s. As the spectrometer imposes a dead time δt of 356 μs during which no signal is acquired, this represents a waste of $356/(356 + 1000) \approx 26\%$ of the total acquisition time. To reduce the total acquisition time, it is possible to reduce the number of patterns $2K$ uploaded onto the DMD. The total acquisition time of the accelerated acquisition depends directly on the number of patterns (e.g., $11.5/2 \approx 5.75$ s considering only half of the patterns). This acceleration comes at the cost of a spatial resolution reduction, as illustrated in Fig. 8. The acceptable acceleration factor depends on the frequency content of the scene. While a 2-fold acceleration may be already excessive for sharp or highly structured objects, an acceleration up to 8-fold may be acceptable for smoother objects. In this study, we chose a subsampling strategy based on energy criteria; however, any other subsampling strategy can be evaluated by subsampling a full acquisition *a posteriori*. The determination of the best subset of patterns remains an open problem and the subject of active research. To ease such studies, we provide several examples for reading and reordering the SPIHIM measurements according to a given subsampling pattern as done in Fig. 8.

For the sake of generality, we have chosen to provide reconstructions where no assumptions are made on the spectral content of the hypercube. Therefore, our algorithm reconstructs each λ -slice of the hypercube independently using a DC-Net. We have chosen this approach for its robustness to noise deviation, as underlined in [21]. However, many other data-driven approaches can be considered, including approaches that exploit the spatio-spectral redundancy to jointly reconstruct several λ -slices (e.g., [46]). The OpenSpyrit ecosystem could serve as a basis for a more systematic comparison of the performance of these algorithms. Alternatively to the splitting strategy given by Eq. (2), the patterns can be shifted to positive values. While the comparison of shifting and splitting have been addressed in [31] for wavelet patterns, the question of handling the negative values of Hadamard patterns remains open. The SPIHIM collection allows studies

to be conducted with both approaches. Indeed, shifted Hadamard patterns $\frac{1}{2}(SH + 1)$ coincide with the positive part of Hadamard patterns P^+ . Therefore, the SPIHIM collection includes the measurements from shifted Hadamard patterns that are obtained by retaining every second spectrum.

6. Conclusion

We propose OpenSpyrit, an ecosystem for reproducible research in hyperspectral single-pixel imaging. In particular, we introduce SPAS (a Python acquisition package), SPYRIT (a Python reconstruction package) and SPIHIM (a data collection), to respond to the need for reproducibility and open access in single-pixel imaging. The SPIHIM collection currently contains 140 hypercubes that are natively $64 \times 64 \times 2048$ in size with a spectral resolution of 2.3 nm and a spatial resolution that can be adjusted between 182.4 μm and 15.2 μm using a digital zoom. It also contains the hypercubes reconstructed at resolution of $128 \times 128 \times 2048$ by a data-driven reconstruction algorithm. The SPIHIM dataset can be exploited for spectral imaging studies in general, but also for more specific studies. In particular, the OpenSpyrit ecosystem should provide a benchmark for single-pixel reconstruction algorithms. In the future, the data collection is expected to continue growing and, in a similar manner, the SPYRIT package should integrate novel reconstruction algorithms.

Funding. Institut Universitaire de France; Institut des sciences de l'information et de leurs interactions (PNRIA ARIM); Agence Nationale de la Recherche (ANR-11-LABX-0063, ANR-17-CE19-0003; ANR-22-CE19-0030-01).

Acknowledgments. Part of this work is based on acquisitions achieved on the PILoT facility, member of the infrastructure France Life Imaging (ANR-11-INBS-0006). N.D. would like to thank Antonio Lorente Mur for the many fruitful discussions on reconstruction strategy and Fadoua Taia-Alaoui who participated in the project as part of the network of engineers of the CNRS National Research Programme in Artificial Intelligence.

Disclosures. The authors declare that they have no conflicts of interest.

Data availability. Data underlying the results presented in this paper are available in the SPIHIM collection of the PILoT warehouse [26].

Supplemental document. See [Supplement 1](#) for supporting content.

References

1. B. Boldrini, W. Kessler, K. Rebner, and R. W. Kessler, "Hyperspectral Imaging: A Review of Best Practice, Performance and Pitfalls for in-line and on-line Applications," *J. Near Infrared Spectrosc.* **20**(5), 483–508 (2012).
2. Q. Li, X. He, Y. Wang, H. Liu, D. Xu, and F. Guo, "Review of spectral imaging technology in biomedical engineering: Achievements and challenges," *J. Biomed. Opt.* **18**(10), 100901 (2013).
3. A. F. H. Goetz, "Three decades of hyperspectral remote sensing of the Earth: A personal view," *Remote. Sens. Environ.* **113**, S5–S16 (2009).
4. J. W. Lichtman and J.-A. Conchello, "Fluorescence microscopy," *Nat. Methods* **2**(12), 910–919 (2005).
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 Process. Mag.* **33**(5), 95–108 (2016).
6. G. R. Arce, D. J. Brady, L. Carin, H. Arguello, and D. S. Kittle, "Compressive Coded Aperture Spectral Imaging: An Introduction," *IEEE Signal Process. Mag.* **31**(1), 105–115 (2014).
7. Y. Oiknine, I. August, V. Farber, D. Gedalin, and A. Stern, "Compressive Sensing Hyperspectral Imaging by Spectral Multiplexing with Liquid Crystal," *J. Imaging* **5**(1), 3 (2018).
8. K. Monakhova, K. Yanny, N. Aggarwal, and L. Waller, "Spectral DiffuserCam: Lensless snapshot hyperspectral imaging with a spectral filter array," *Optica* **7**(10), 1298 (2020).
9. T. Sun and K. Kelly, "Compressive Sensing Hyperspectral Imager," in *Frontiers in Optics 2009/Laser Science XXV/Fall 2009 OSA Optics & Photonics Technical Digest*, (OSA, 2009), p. CTuA5.
10. F. Soldevila, E. Irlés, V. Durán, P. Clemente, M. Fernández-Alonso, E. Tajahuerce, and J. Lancis, "Single-pixel polarimetric imaging spectrometer by compressive sensing," *Appl. Phys. B* **113**(4), 551–558 (2013).
11. J. Peller, F. Farahi, and S. R. Trammell, "Hyperspectral imaging system based on a single-pixel camera design for detecting differences in tissue properties," *Appl. Opt.* **57**(27), 7651–7658 (2018).
12. F. Rousset, N. Ducros, F. Peyrin, G. Valentini, C. D'Andrea, and A. Farina, "Time-resolved multispectral imaging based on an adaptive single-pixel camera," *Opt. Express* **26**(8), 10550–10558 (2018).
13. M. P. Edgar, G. M. Gibson, and M. J. Padgett, "Principles and prospects for single-pixel imaging," *Nat. Photonics* **13**(1), 13–20 (2019).

14. G. M. Gibson, S. D. Johnson, and M. J. Padgett, "Single-pixel imaging 12 years on: A review," *Opt. Express* **28**(19), 28190–28208 (2020).
15. M. Harwit and N. J. A. Sloane, *Hadamard Transform Optics* (Academic Press, 1979).
16. E. D. Nelson and M. L. Fredman, "Hadamard Spectroscopy," *J. Opt. Soc. Am.* **60**(12), 1664–1669 (1970).
17. E. J. Candes and M. B. Wakin, "An Introduction To Compressive Sampling," *IEEE Signal Process. Mag.* **25**(2), 21–30 (2008).
18. G. Barbastathis, A. Ozcan, and G. Situ, "On the use of deep learning for computational imaging," *Optica* **6**(8), 921–943 (2019).
19. C. F. Higham, R. Murray-Smith, M. J. Padgett, and M. P. Edgar, "Deep learning for real-time single-pixel video," *Sci. Rep.* **8**(1), 2369 (2018).
20. S. Rizvi, J. Cao, K. Zhang, and Q. Hao, "DeepGhost: Real-time computational ghost imaging via deep learning," *Sci. Rep.* **10**(1), 11400 (2020).
21. A. Lorente Mur, P. Leclerc, F. Peyrin, and N. Ducros, "Single-pixel image reconstruction from experimental data using neural networks," *Opt. Express* **29**(11), 17097–17110 (2021).
22. A. Lorente Mur, F. Peyrin, and N. Ducros, "Deep Expectation-Maximization for Single-Pixel Image Reconstruction With Signal-Dependent Noise," *IEEE Trans. Comput. Imaging* **8**, 759–769 (2022).
23. F. Wang, C. Wang, C. Deng, S. Han, and G. Situ, "Single-pixel imaging using physics enhanced deep learning," *Photonics Res.* **10**(1), 104 (2022).
24. Y. Tian, Y. Fu, and J. Zhang, "Plug-and-play algorithms for single-pixel imaging," *Opt. Lasers Eng.* **154**, 106970 (2022).
25. Y. Lu, B. Tan, S. Ding, and Y. Li, "SPI-CGAN: Single-pixel imaging based on conditional generative adversarial network," *J. Phys.: Conf. Ser.* **2216**(1), 012092 (2022).
26. L. Mahieu-Williame and N. Ducros, "Single-Pixel Hyperspectral Image (SPIHIM) Collection," Pilot Warehouse, 2022, <https://pilot-warehouse.creatis.insa-lyon.fr/#collection/6140ba6929e3fc10d47dbe3e>.
27. M. D. Wilkinson, M. Dumontier, and I. J. Aalbersberg, *et al.*, "The FAIR Guiding Principles for scientific data management and stewardship," *Sci. Data* **3**(1), 160018 (2016).
28. G. Beneti Martins, L. Mahieu-Williame, and N. Ducros, "Single-pixel acquisition software," (version 1.1), Github, 2023, <https://github.com/openspyrit/spas>
29. A. Lorente Mur, T. Baudier, and N. Ducros, "Single-pixel reconstruction toolkit," (version 2.0.0), *GitHub*, (2022), <https://github.com/openspyrit/spyrit>
30. E. M. V. Association, "EMVA standard 1288, standard for characterization of image sensors and cameras," Release **3**, 29 (2016).
31. A. Lorente Mur, M. Ochoa, J. Cohen, X. Intes, and N. Ducros, "Handling negative patterns for fast single-pixel lifetime imaging," in *Molecular-Guided Surgery: Molecules, Devices, and Applications V*, B. W. Pogue and S. Gioux, eds. (SPIE, 2019), p. 9.
32. P. G. Vaz, D. Amaral, L. F. R. Ferreira, M. Morgado, and J. Cardoso, "Image quality of compressive single-pixel imaging using different Hadamard orderings," *Opt. Express* **28**(8), 11666–11681 (2020).
33. A. Lorente Mur, P. Bataille, F. Peyrin, and N. Ducros, "Deep Expectation-Maximization For Image Reconstruction From Under-Sampled Poisson Data," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, (IEEE, 2021), pp. 1535–1539.
34. S. M. Popoff and M. W. Matthès, "ALP4lib: A Python wrapper for the Vialux ALP-4 controller suite to control DMDs," Zenodo, 2020, <https://zenodo.org/record/4076193#ZCWKGnbMJdg>.
35. Measurement Standards Laboratory of New Zealand, "MSL-Equipment,".
36. H. S. Fairman, M. H. Brill, and H. Hemmendinger, "How the CIE 1931 color-matching functions were derived from Wright-Guild data," *Color Res. Appl.* **22**(1), 11–23 (1997).
37. A. Paszke, S. Gross, and F. Massa, *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, eds. (Curran Associates, Inc., 2019), pp. 8024–8035.
38. N. Ducros, A. Lorente Mur, and T. Baudier, "Spyrit-examples," (version 2.0.0), *GitHub*, 2023, <https://github.com/openspyrit/spyrit-examples/>.
39. A. Coates, A. Ng, and H. Lee, "An Analysis of Single-Layer Networks in Unsupervised Feature Learning," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, (JMLR Workshop and Conference Proceedings, 2011), pp. 215–223.
40. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, (2009), pp. 248–255.
41. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, eds. (2015).
42. A. Orych, "Review of Methods for Determining the Spatial Resolution of UAV Sensors," *The Int. Arch. Photogramm. Remote. Sens. Spatial Inf. Sci.* **XL-1/W4**, 391–395 (2015).
43. L. Baldassarre, Y.-H. Li, J. Scarlett, B. Gözcü, I. Bogunovic, and V. Cevher, "Learning-Based Compressive Subsampling," *IEEE J. Sel. Top. Signal Process.* **10**(4), 809–822 (2016).
44. "AVIRIS - Airborne Visible / Infrared Imaging Spectrometer - Data," <https://aviris.jpl.nasa.gov/data/index.html>.

45. B. Arad and O. Ben-Shahar, "Sparse Recovery of Hyperspectral Signal from Natural RGB Images," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, eds. (Springer International Publishing, 2016), Lecture Notes in Computer Science, pp. 19–34.
46. V. Pronina, A. Lorente Mur, J. F. P. J. Abascal, F. Peyrin, D. V. Dylov, and N. Ducros, "3D denoised completion network for deep single-pixel reconstruction of hyperspectral images," *Opt. Express* **29**(24), 39559 (2021).