

Statistical compressive sensing method for Hadamard-based single-pixel microscopy supported by kernel density estimators

H. Tobon-Maya,^{1*} S. I. Zapata-Valencia,¹ M. Obando,² F. Lucka,² E. Tajahuerce,¹ and J. Lancis¹

¹Institute of New Imaging Technologies (INIT), Universitat Jaume I, Castelló, Spain

²Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

Abstract. Hadamard-based single-pixel microscopy (HSPM) is a versatile non-conventional imaging technique where a binary function base is projected over the sample in a microscope setup to recover its information. One HSPM's main challenge is the need to project numerous patterns to retrieve the image of the object under study. This leads to potential phototoxicity damage and a reduction in temporal resolution. Aiming to reduce the total pattern projection time, this study explores the use of statistical compressive sensing (CS) using the kernel density estimator (KDE) approach to learn the probability distribution of the most relevant Hadamard spectrum (HS) sampling coefficients, based on a large-scale dataset of 50,000 histopathology images. The probability distribution can then be sampled to generate the set of Hadamard patterns to be projected. The proposed KDE-guided CS method is implemented and tested on biological and non-biological samples. An image resolution of 550 lp/mm was recovered at a 25% sampling ratio (SR) using the proposed method, a level not reached by the well-established TV-based approach. Moreover, compared to TV-based sampling, the Michelson contrast increased from 0.09 to 0.17 at a 25% SR and from 0.12 to 0.29 at a 30% SR. These results demonstrate the feasibility of the proposed method for HSPM CS applications.

Keywords: single pixel; microscopy; computational imaging.

Received Sep. 4, 2025; revised manuscript received Oct. 17, 2025; accepted Oct. 31, 2025; published online Dec. 8, 2025.

© The Authors. Published by Chinese Laser Press under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.

[DOI: [10.3788/AI.2025.10001](https://doi.org/10.3788/AI.2025.10001)]

1. Introduction

The use of high-speed spatial light modulators (SLMs) and non-spatially resolved detectors, with broad spectral sensitivity, has allowed the implementation of diverse applications of single-pixel imaging (SPI) techniques^[1] such as structured illumination single-pixel microscopy (SPM)^[2–4]. Unlike conventional imaging systems, in SPM, a scanning base is projected over the sample using an SLM. This process is often performed by employing high-speed digital micromirror devices (DMDs), due to their achievable frame rates, up to 33 kHz, and their wide spectral range. This configuration minimizes the sample-light interaction by constantly changing the illumination over the sample for each pattern of the scanning base, thereby reducing photobleaching and phototoxicity. The technique's adaptability has

attracted attention to different applications, including polarization state^[5] studies, phase imaging^[6], fluorescent lifetime^[6], and imaging outside the visible spectrum^[7], covering wavelengths from terahertz to X-ray regions^[5,8].

Among the multiple alternatives for SPM sampling, deterministic approaches like Hadamard-based and Fourier-based SPM have been broadly implemented^[9]. Its performance under noise and low-intensity sensing conditions suggests the Hadamard approach as a suitable implementation for microscopy techniques, where the signals are often low in intensity and sensitive to environmental variables such as mechanical vibration. In Hadamard-based single-pixel microscopy (HSPM), each pattern of the Hadamard base, corresponding to a two-dimensional rectangular function, is projected over the sample, as shown in Fig. 1(a). Then, the pointwise product between each demagnified Hadamard pattern $H_k(x/M, y/M)$ by the

*Address all correspondence to H. Tobon-Maya, tobon@uji.es

microscope setup and the sample $s(x, y)$ is integrated by a bucket detector. This process allows the measurement of an intensity vector $I_v(k)$, as illustrated in Fig. 1(b). In that way, one can write $I_v = \mathcal{H}(s)$, with $\mathcal{H}(\cdot)$ being the Hadamard transformation. The collected intensity vector contains the Hadamard frequency spectrum (HS) of the imaged sample, which can be reshaped using the Walsh order or the 2D representation, as seen in Fig. 1(c)^[10]. The coordinates f_x, f_y are determined by the frequency of the rectangular waves that compose the Hadamard base. The maximum available frequency for each axis represents a rectangular wave with a semi-period of $1/n_p$ units, where n_p is the number of pixels in each base element along each axis. Finally, when the full spectrum is recovered, the sample image can be obtained just by applying the inverse Hadamard transform $\mathcal{H}^{-1}(\cdot)$ as shown in Fig. 1(d).

One significant trade-off present in this imaging technique is the need to project a large number of binary patterns to recover accurate information about the object under study. Recently, this improvement opportunity has been a topic of interest, leading to different innovative approaches for compressive sensing (CS) in SPI and HSPM^[11–14]. Most of the methods are based on a proper sampling sequence of Hadamard patterns, whereby the more relevant patterns are scanned first and those providing low information are avoided^[11,14,15]. Among them, the total variation order or ZigZag order (TV/ZigZag) is probably the most intuitive and the one with the best performance to apply CS approaches for SPI^[12,14,15]. In the 2D Walsh representation, shown in Fig. 1(c), this CS sampling method represents the analog to measure following a low-pass triangular filter over the top-left corner of the HS^[16]. Nevertheless, as will be shown later in this study, the relevant image information in the HS is often distributed among both the low- and the high-frequency domains.

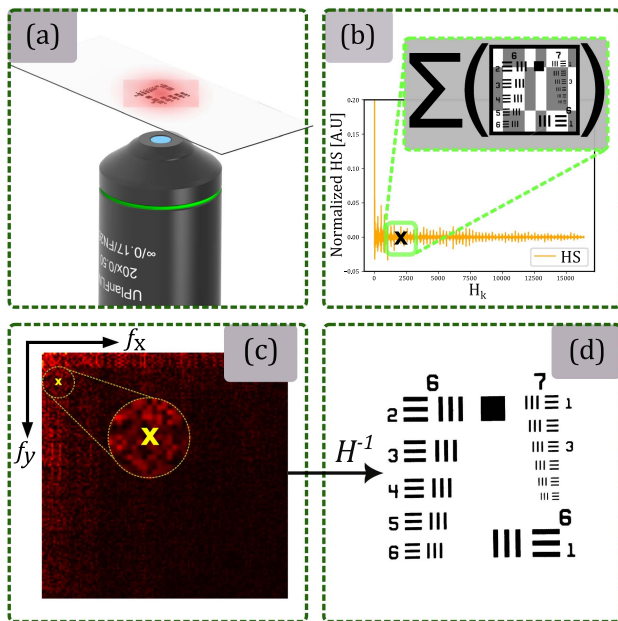


Fig. 1 HSPM schematic. Panel (a) shows the demagnified Hadamard patterns over the sample. Panel (b) illustrates the intensity vector construction. Panels (c) and (d) represent the resulting Hadamard transform in the 2D Walsh representation and the final image after the \mathcal{H}^{-1} implementation, respectively.

As an alternative to deterministic sampling approaches, recent alternatives involve random sampling based on small dataset analysis, which introduces probabilistic elements into the sampling process^[16]. This method considers both low- and high-frequency components. Instead of being conditioned by a triangular filter, the sampling will be conditioned on a pre-determined exponential model for the analyzed HS energy in the studied dataset, where the probability decays in terms of the Euclidean distance to the zero-frequency order^[16]. As well, some learning-based SPI has been recently proposed in the context of SPI increasingly relying on learning-based pipelines^[17–20], spanning end-to-end neural decoders and learned/optimized sampling. While these approaches can be effective, they often depend on sizeable training sets, tuned hyperparameters, and, in several cases, non-public code or weights, which complicates like-for-like experimental comparison and reproducibility. With the aim of contributing to a more adaptable probabilistic sampling method specifically designed for HSPM, this study explores the statistical behavior of HS to propose a probabilistic sampling kernel for CS in HSPM. By analyzing the HS of 50,000 images through kernel density estimation (KDE), this study seeks to define a sampling approach that adjusts to diverse imaging conditions. Such a method could provide an alternative to current deterministic and probabilistic approaches, enabling more efficient and adaptable SPI in both low- and high-frequency domains.

2. Hadamard Spectrum Information Distribution

The broad implementation of CS in Hadamard-based SPI is primarily supported by the natural sparsity of images in the HS. When CS is applied, only a portion of the total HS is sampled, leading to an underdetermined system that is often solved using algorithms such as TVAL3^[21] or NESTA^[22]. This natural sparsity is demonstrated for 4 different images, shown in Figs. 2(a)–2(d). The magnitude of their corresponding normalized HS coefficients is displayed in Figs. 2(e)–2(h), illustrating that most HS coefficients do not encode significant sample information. The approximated compressive ratio, ACR, of each sample is a commonly used, ad hoc way to quantify the proportion of significant HS coefficients and is defined as

$$\text{ACR} = \frac{\sum_{k=1}^T \mathbf{1}(|H_k(s)| > \text{mean}_{|H(s)|})}{T}, \quad (1)$$

where $H(s)$ represents the Hadamard transform of an arbitrary image $s(x, y)$, and T is the total number of pixels in the image. The ACR, thus, determines that all HS coefficients that are larger than the mean encode the most relevant sample information. For each image, the coefficients contributing to the ACR are shown in Figs. 2(i)–2(l).

These coefficients define a subsampling base for the HS. This subsampling base will be referred to from now on as the ACR filter. In Figs. 2(m)–2(p), images of the reconstruction achieved by applying \mathcal{H}^{-1} to the selected coefficients using the ACR filter are shown. This yields images with a structural similarity index measure (SSIM)^[17] from 0.84 to 0.91 and a signal-to-noise ratio (SNR), ranging from 70 to 77 dB, as indicated in the corresponding reconstruction panels. Both metrics were computed as defined in the metrics section of the Python scikit-image package described in Ref. [23]. Importantly, this




Fig. 2 Hadamard SPI behavior. In the first column, panels (a) to (d), the images of a church tower, a parrot, a flower, and a histological microscopy sample are shown, respectively. The corresponding Hadamard transform of each image is displayed in panels (e) to (h). The ACR mask filter for each HS is shown in panels (i) to (l), respectively. The image reconstruction using the ACR filter without any CS algorithm is shown in panels (m) to (p), respectively.

reconstruction does not involve solving the CS inverse problem through any optimization algorithm. Instead, the non-sampled coefficients are set to zero. However, this approach requires prior knowledge of the full HS of the imaged sample. Each ACR filter encompasses sampling ratios (SRs) around 0.21 of the total data. One can validate that the corresponding TV/Zig-Zag sampling filter with an SR of 0.21 only encodes around 52% of the relevant image coefficients, leaving behind relevant data for a full CS data acquisition scenario.

This study proposes an alternative to the above scenario by analyzing different ACR filters to define a general sampling kernel to be used as a measurement input for CS in HSPM. The similarity in the behavior of the distributions of ACR filters in Fig. 2 motivates the search for a function to generalize the sampling of relevant HS coefficients. To address this, we propose to use a probabilistic function for defining a sampling kernel for HSPM. This function is constructed using the KDE algorithm, an unsupervised machine learning method. To generalize the approach, the analysis will be extended to a large dataset of images, enabling the generation of robust sampling kernels.

3. Methods

3.1. Kernel Density Estimator

Determining the probability density function (PDF) to adjust the ACR filter data presented in Fig. 2 is a challenging task. Standard fitting approaches rely on predefined functions, such as gamma or exponential distributions. However, these methods assume that the data follows a specific behavior, which may not be accurate for this application. The ACR filter distribution, especially for high-frequency components, can be irregular and sample-dependent. This irregularity becomes evident in the superposed data presented in Figs. 2(i) to 2(l), where the combined ACR filters show a distribution with a non-conventional behavior. The high-frequency components exhibit variations that depend on the specific imaged sample, further complicating the analysis. In such cases, the KDE provides an effective and flexible method for estimating the underlying data distribution. Unlike parametric approaches, KDE does not assume a specific functional form for the global PDF. Instead, it estimates the density by placing local kernels, typically Gaussian, at each data point and summing their contributions^[24,25]. KDE has been

successfully applied in various complex scenarios^[26-28], making it an ideal choice for analyzing the distributions of ACR filters proposed in this study. Given the normalized frequency coordinates $u = f_x/f_{cx}$, $v = f_y/f_{cy}$ for the ACR filter, with f_{cx} and f_{cy} being the corresponding cutting frequency along each axis, the KDE method can be used to estimate the PDF. In one dimension, the KDE is mathematically defined as

$$\hat{f}(u) = \frac{1}{n(w_u)^d} \sum_{i=1}^n K\left(\frac{u - u_i}{w_u}\right), \quad (2)$$

where $\hat{f}(u)$ is the estimated density at point u , n is the total number of data points, w_u is the bandwidth parameter controlling the kernel's width, n is the dimensionality of the data, and $K(\cdot)$ is the kernel function, commonly chosen as a Gaussian, $K(g) = \frac{1}{\sqrt{2\pi}} e^{-\frac{g^2}{2}}$. For a 2D analysis, the KDE extends naturally as

$$\hat{f}(u, v) = \frac{1}{nw_u w_v} \sum_{i=1}^n K\left(\frac{u - u_i}{w_u}\right) K\left(\frac{v - v_i}{w_v}\right), \quad (3)$$

where w_u and w_v are the bandwidths along the u and v axes, respectively. To implement this in our analysis, the relative positions (u, v) of the ACR-relevant coefficients are used as input. The KDE then evaluates the local density of these points, creating a continuous PDF representing the spatial distribution, having each HS coefficient's occurrence as a contribution. In that way, by picking the dataset relevant sampling f_x, f_y coefficients based on the ACR filter, the local PDF composed by summing different bandwidths around each observation $P(|H_k(f_x, f_y)| > \text{mean}(|H(f_x, f_y)|))$ can be computed using the KDE. This method provides an accurate PDF estimation for irregular data, while its reliability depends on the quantity of data input. A small dataset might not capture the full range of variability in the HS coefficients.

3.2. KDE Sampler Using a Large Microscopy Image Dataset Analysis

To ensure a robust estimation, this study builds a comprehensive dataset of relevant coefficients based on a large dataset analysis^[29,30]. A total of 50,000 images of histopathologic scans of lymph node sections extracted from the PatchCamelyon dataset^[29,30] were used. The dataset is composed of histopathology images exhibiting a variety of spatial frequency content. The diverse frequency content minimizes the introduction of non-desired prior information in the frequency domain and ensures variability in the distribution. This approach ensures that the sampling kernel reflects a generalizable and accurate representation of the data, enabling consistent performance for HSPM observations. Figure 3 describes the general process of data construction to generate the inputs for the KDE fitting and the model training. As illustrated in the flow chart in Fig. 3, first a random sample of 50,000 images of the original data set was taken and upsampled from 96 pixel \times 96 pixel to 128 pixel \times 128 pixel, and then the Hadamard transformation was computed. Later, the ACR filter was computed for all the images, and the positions for the relevant coefficients were condensed as the input for the KDE model. All the frequencies were normalized with respect to the cutting frequencies f_c , so, in that way, each data point to be adjusted will return a relative index coefficient to the implemented HS order.

Once the dataset was created, the KDE model was fitted over the data using the Google Cloud Platform (GCP) facilities. The Python packages provided in the scikit-learn for KDE adjustment were implemented in the construction of our ACR sampling filter^[31,32]. The model fitting was performed in a 16-core virtual machine with 64 GB of RAM. The partial fitting method was implemented to generate different samplers depending on the final user's computing capabilities. All the codes regarding the use of the KDE samples and generation of the imaging patterns can be found in the data availability section referenced repository.

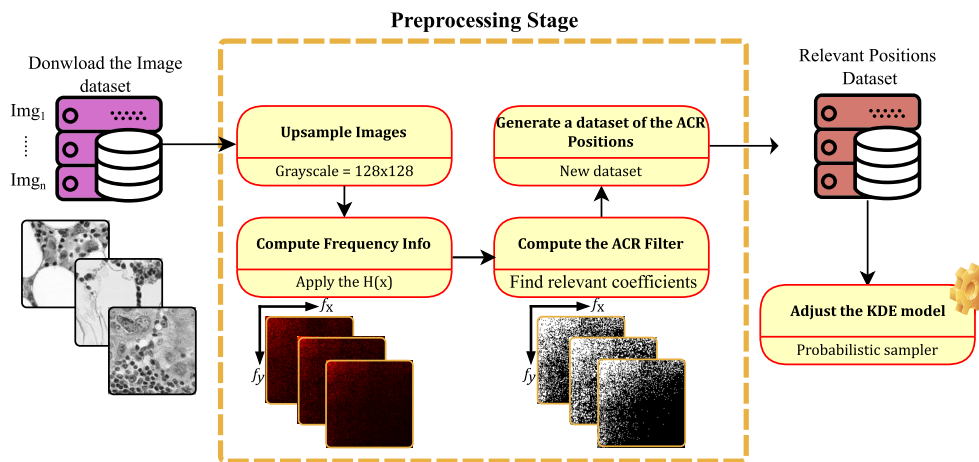


Fig. 3 Data construction flow chart for KDE fitting. First, 50,000 histopathology images were randomly sampled from the PatchCamelyon dataset and upsampled from 96 pixel \times 96 pixel to 128 pixel \times 128 pixel. Next, the Hadamard transformation was computed on each image. The ACR filter was then computed to identify the most informative coefficient positions. All frequencies were normalized with respect to the cutoff frequency, yielding relative index coefficients. Finally, the condensed coefficient positions were supplied as input to the kernel density estimator to generate the probabilistic sampling kernel.

3.3. KDE Sampler Performance

The KDE sampler can be used to generate different ACR filters according to the desired SR. The model global performance to determine the sampling coefficients for an image that was not part of the histopathologic fitting data is shown in Fig. 4. A simulated Siemens star image of 128 pixel \times 128 pixel, shown in Fig. 4(a), was used to validate the sampler performance encoding the ACR filter coefficients. Since the KDE returns a probability function, the sampling filters will be slightly different every time they are generated. To begin with the model performance determination, an SR of 25% was chosen, and different filters were computed recursively. In each iteration, the number of coincidences between the original ACR filter of the Siemens star, shown in Fig. 4(b), and the corresponding KDE-generated filters, one of which is shown in Fig. 4(c), were stored. A histogram of the coincidences alongside 500 iterations of this correspondence comparison is shown in Fig. 4(d). This process was performed with KDE filters computed with 0.5, 0.8, and 1.1 bandwidths. The mean value of correspondence falls around 0.72 for all the trained distributions, as shown in Fig. 4(c).

As the corresponding distributions in Fig. 4(d) are essentially indistinguishable across the three bandwidth candidates $w \in \{0.5, 0.8, 1.1\}$, with similar mean, variance, and overall shape. We adopt the intermediate bandwidth equal to 0.8 for all subsequent performance tests. The imaging performance evaluation of the proposed method under a simulated scenario using different SRs is shown in Fig. 5. To conduct the test, SRs of 20%, 25%, and 30% are evaluated. To retrieve the final

images in this study, both in the simulated and experimentally acquired scenarios, the CS inverse problem, to retrieve the final image, is addressed by implementing the well-known NESTA optimization routine^[22]. A detailed explanation of the CS reconstruction is provided in the [Supplement 1](#), including the inverse problem formulation and the NESTA-based minimization procedure.

The proposed method is compared with the well-known TV/Zigzag method^[12,14,15] for imaging the simulated Siemens star test, as shown in Figs. 5(a)–5(c). The KDE performance is shown in Figs. 5(d) and 5(e) for the same sample. Despite being able to retrieve high-frequency information, unlike the TV/ZigZag method, background noise is persistent in the retrieved images. This behavior can be related to non-uniformity in the sampled HS low-frequency components. To compensate for this situation, a 5% SR low-pass filter, based on the TV/ZigZag sampling, is added to the KDE filter. In that way, we ensure a homogeneous sampling of the background as seen in Figs. 5(f) and 5(g). The KDE + low frequency form TV/ZigZag does not imply any change in the problem solution. This augmentation does not change the optimization problem or the solver; it simply adds a few low-frequency measurements before reconstruction, which improves background uniformity while preserving the high-frequency benefits of KDE.

The capability of resolving small details in different objects is crucial in microscopy and represents one of the biggest trade-offs for the integration of CS approaches in SPM. One notable advantage of the proposed method, compared with the well-known TV sampling technique, is the resolution impact. This

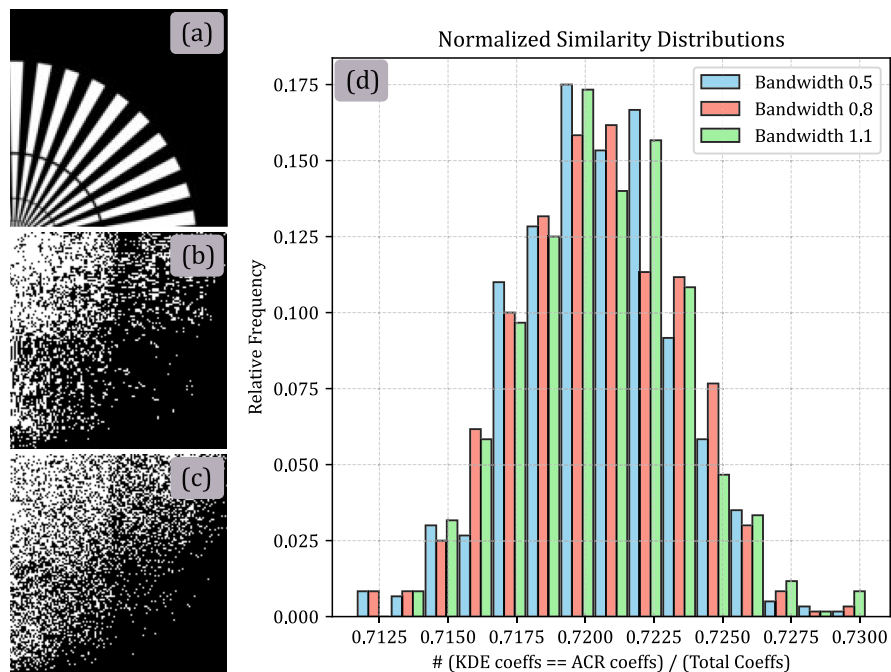


Fig. 4 KDE sampler performance for an image that is distinct from the histopathologic training data. A Siemens star test and its corresponding HS ACR are shown in panels (a) and (b), respectively. In panel (c), a kernel for a 25% SR is shown. In panel (d), a histogram plotting the coincidences between 500 generated KDE sampling filters and the Siemens star HS ACR is shown, demonstrating the method's capability of producing a sampling kernel with up to 72% of coincidence.

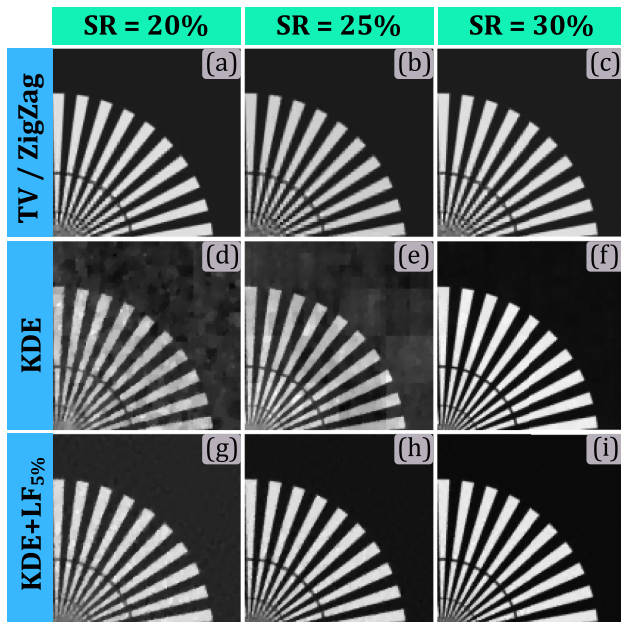


Fig. 5 Simulated reconstruction performance on a 128×128 Siemens star test target under three SRs. (a)–(c) Reconstructions obtained with TV/ZigZag sampling at SR = 20%, 25%, and 30%, respectively. (d)–(f) Reconstructions obtained with the KDE-guided sampler at the same SRs, showing improved high-frequency fidelity but nonuniform background sampling. (g)–(i) Reconstructions incorporating a 5% low-pass (LP) complement on the KDE kernel, and the filter restores uniform background sampling while preserving high-frequency detail.

can be noticed in Fig. 5, where the inner spikes of the Siemens star test are better resolved by implementing the proposed method. This behavior is present as well in the reconstruction of experimental samples, as will be shown in the Sec. 4. Additionally, to provide a fair, reproducible context, the Supplement 1 reports the SSIM and SNR against transparent baselines, such as random Hadamard under sequency and conventional orderings and scrambled-Hadamard, across 15%–60% SRs. At equivalent quality, the sampling burden is

substantially reduced: to attain $SSIM \approx 0.95$, most of the traditional sampling approaches required $\sim 50\%$ SR, while the proposed method achieves this at $\sim 30\%$ SR, as shown in detail for different methods in the Supplement 1.

4. Experimental Results

4.1. Experimental Setup

The experimental validation of the KDE-guided sampling approach was performed on a custom telecentric single-pixel microscope, as schematized in Fig. 6. The microscope was built following the procedure reported in Ref. [2]. The core imaging train combines a Mitutoyo 20 \times infinity-corrected microscope objective (MO, NA=0.42) with a 200 mm tube lens (TL). Downstream of the TL, a 4 f relay (L1 and L2, both $f = 75$ mm) provides access to the system’s pupil plane. An Optotune ETL-12-30 tunable lens with a range of ± 5 diopters is inserted to introduce focus compensation.

A Vialux V-6005 DMD is conjugated both to the sample plane and to the microscope field stop, ensuring that each projected Hadamard pattern is sharply demagnified onto the specimen. Illumination is provided by a fiber-coupled 650 nm diode. The light beam is collimated, expanded to fill the DMD active area, and then reflected through the objective onto the sample. Reflected light by the sample retraces the optical path, reaches a 50:50 beam splitter (BS), and is collected by a second MO into an avalanche photodiode APD410A/M, acting as the bucket detector shown in Fig. 6. The photodiode output is digitized by a 16-bit ADC at 200 kHz in synchrony with the DMD using a NI-DAQ USB 6341. Finally, the pattern projection rate at the DMD is fixed at 9 kHz.

4.2. Experimental Performance

To quantitatively assess the imaging capabilities of the proposed KDE-guided sampling strategy, a resolution test chart and a cotton-tissue sample were imaged. All the experiments were performed using the described setup. Figure 7 presents the results obtained on an Edmund Optics resolution test chart, wherein the KDE-guided method is compared against the established TV/ZigZag approach at SRs of 25% and 30%.

In Figs. 7(a) and 7(b), line pairs up to 550 lp/mm were resolved using the KDE-guided strategy at both 25% and 30%




Fig. 6 Experimental HSPM validation setup schematic. A 650 nm collimated source illuminates a Vialux V-6005 DMD. The DMD was relayed through a 4 f system (L1, L2; $f = 75$ mm), conjugated to both the sample plane and the microscope field stop. An Optotune ETL-12-30 tunable lens (TL) was inserted within this relay to enable rapid focus adjustment. Patterned illumination is projected through a 20 \times , NA 0.42 microscope objective (MO), and a 200 mm TL onto the sample. Backscattered or fluorescent light was separated by a dichroic/beam splitter (BS), collected by a second objective (MO₂), and detected by an avalanche photodiode (APD).

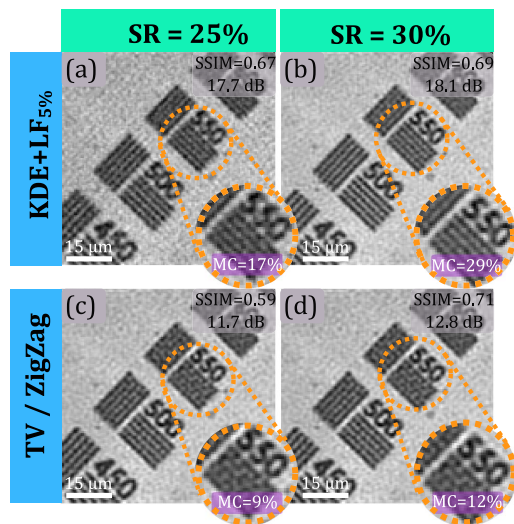


Fig. 7 Experimental validation of the proposed method using a high-resolution test target. In panels (a) and (b), the proposed method retrieved images KDE + TV retrieved images for the test target at 25% and 30% SRs are shown, respectively. In panels (c) and (d), the TV/ZigZag retrieved images for the test target at 25% and 30% SRs are shown, respectively. The scale bars for all the images correspond to 15 μm .

SRs. These features remained unresolvable with TV/ZigZag with the same SR, shown in Figs. 7(c) and 7(d). Within the 550 lp/mm region, highlighted by the orange insets, the Michelson contrast (MC) was measured to provide a quantitative metric of the method performance. The MC is shown in rectangular purple insets for each image. Contrast values of 17% and 29% were measured for the KDE-guided method in the zoomed insets in Figs. 7(a) and 7(b), respectively. Under identical conditions, TV/ZigZag yielded contrasts for the zoomed inset of 9% and 12%, respectively, corresponding to a twofold contrast enhancement at 25% SR and a 2.4 \times improvement at 30% SR.

Moreover, at 500 lp/mm, resolution was achieved by both methods; however, a higher contrast of 30% was measured with KDE guidance, compared to 24% for TV/ZigZag at 25% SR. These results demonstrate that, even under high compression rates, the proposed approach preserves high-frequency detail while substantially enhancing contrast in the finest features, both for high and intermediate frequency values. Additionally, the SSIM and SNR were evaluated relative to a full-sampling reference. At 20% SR, the TV/ZigZag baseline attains a 12.8 dB SNR, whereas the proposed KDE-guided sampling achieves 18.1 dB under identical conditions. A detailed analysis across SRs and sampling policies for the resolution target is provided in the [Supplement 1](#).

The performance of the proposed KDE-guided sampling strategy was also evaluated on a biological specimen by imaging the fluorescence emission of a cotton-tissue sample (Fig. 8). Excitation was provided by a 450 nm diode laser, and fluorescence was separated from the illumination path by replacing the 50:50 BS (Fig. 6) with a Thorlabs MD498 dichroic mirror. Focus alignment was ensured via the Optotune FTL, operated according to the procedure described in Ref. [33].

Three representative reconstructions of the fluorescent cotton-tissue sample are presented side by side in Fig. 8. In




Fig. 8 Fluorescence reconstructions of a cotton tissue sample acquired with full and compressed Hadamard sampling. In panel (a), the reference image obtained at 100% SR is shown. In (b), reconstruction using TV/ZigZag sampling at 25% SR is shown. In (c), the reconstruction using KDE-guided sampling at 25% SR is presented. The scale bars correspond to 20 μm .

Fig. 8(a), the full Hadamard base yields a high-fidelity image in which individual cellulose fibers and their branching points are crisply delineated. The background fluorescence appears as well uniformly low. In Fig. 8(b), the TV/ZigZag, at 25% SR, is shown; many of the fine fibrous details observed in Fig. 8(a) are obscured. Sharp fiber–fiber intersections are blurred, and the contrast between the fibers and the surrounding matrix is substantially reduced, as highlighted by the lower-right orange inset. By contrast, in Fig. 8(c), the KDE-guided sampling, at 25% SR, recovers both the thin fiber strands and the localized bright spots with clarity nearly matching the full-scan reference. Moreover, the background noise floor remains suppressed, resulting in an overall contrast enhancement that closely approaches that of the 100% SR acquisition. The scale bars in all panels correspond to 20 μm . These results illustrate that, even at one-quarter of the full sampling load, the proposed method can preserve high-resolution structural information and achieve superior signal-to-background distinction in fluorescence imaging of biological tissues.

Taken together, our experiments demonstrate that the proposed KDE-guided sampling is the way to high-fidelity final reconstructions at substantially reduced SRs, consistently outperforming TV/ZigZag and random Hadamard orderings across both resolution targets and biological specimens. Extensive assays, including SSIM/PSNR versus SR curves, are provided in the [Supplement 1](#) for all the presented results. In contrast to learning-based SPI approaches, our method attains competitive quality without training or heavy computing, remains decoder-free (standard CS reconstruction via NESTA), and is fully open and reproducible. These attributes, combined with the observed performance, make the proposed sampling method a practical approach for deployment on standard SPI hardware.

5. Conclusion

HSPM has been presented as a versatile, non-conventional imaging modality in which binary Hadamard patterns are projected onto a sample to recover spatial information. The principal limitation of HSPM, namely the requirement to display a large number of patterns, which could induce phototoxicity or photobleaching and lower the achievable framerate for imaging dynamic processes in samples, has been addressed by introducing a statistical CS framework for reducing the needed number of patterns, maintaining resolution and contrast.

A large-scale dataset of 50,000 histological images has been analyzed to identify the most informative Hadamard-spectrum coefficients for each image. This coefficient distribution set has been processed via a KDE to yield a probabilistic sampling function that chooses the most likely coefficients to provide relevant sample information. The KDE-guided CS method has been experimentally validated on a telecentric single-pixel microscope equipped with a DMD, a focus-tunable lens, and NI-DAQ readout, without any hardware modification.

Experimental validations have been performed on both resolution test charts and biological specimens. In resolution chart experiments, line pairs up to 550 lp/mm have been recovered at 25% and 30% SRs, whereas conventional TV/ZigZag sampling fails to resolve these features. The MC has been increased from 0.09 to 0.17 at 25% SR and from 0.12 to 0.29 at 30% SR, representing up to a 2.4× enhancement in the finest details. In fluorescence imaging of cotton tissue at 25% SR, fibrous structures have been preserved with clarity approaching full-scan acquisitions, while traditional CS exhibits blurred features and reduced contrast.

The proposed KDE-guided policy has achieved a given image quality at substantially lower SRs than the TV/ZigZag baseline. For the Siemens-star target, reaching SSIM \approx 0.95 requires \sim 50% SR with TV/ZigZag, whereas the proposed method attains the same quality at \sim 30% SR (see the [Supplement 1](#)). This reduction directly translates into fewer projected patterns, faster acquisition, and lower optical dose.

These results demonstrate that KDE-guided sampling can effectively reduce the number of pattern projections, mitigate photodamage risk, and maintain high spatial resolution and contrast under aggressive compression. The openness of the model facilitates reproducibility and integration on standard hardware. Future work will explore block-adaptive/feedback sampling, video-rate operation, and hardware acceleration while preserving the decoder-free reconstruction pipeline. The method has been shown to be feasible and advantageous for HSPM CS applications in both physical and biological contexts.

Data Availability

The data and code related to the article are openly available on Zenodo at <https://doi.org/10.5281/zenodo.17373219>. The repository includes trained KDE sampling models, deterministic TV/ZigZag sampling kernels in the Walsh–Hadamard domain, and NESTA-based reconstruction scripts.

Acknowledgments

This work was supported by the European Union (No. GA 101072354). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. We also acknowledge the grant (No. PID2022-142907OB-I00) funded by MICIU/AEI/10.13039/501100011033 and “ERDF/EU” and the grant (No. CIPROM/2023/44) funded by the Generalitat Valenciana. This material is based upon work supported by the Google Cloud Research Credits program.

References

1. M. P. Edgar, G. M. Gibson, and M. J. Padgett, “Principles and prospects for single-pixel imaging,” *Nat. Photonics* **13**, 13 (2019).
2. S. I. Zapata-Valencia *et al.*, “Design, assembly, alignment and application of a versatile, open-source, single-pixel microscope,” *Sci. Rep.* **15**, 17802 (2025).
3. A. Ebner *et al.*, “Diffraction-limited hyperspectral mid-infrared single-pixel microscopy,” *Sci. Rep.* **13**, 281 (2023).
4. Y.-N. Zhao *et al.*, “Single-pixel dual-mode microscopy for simultaneous acquisition of magnitude and wrapped phase images,” *Opt. Laser Technol.* **182**, 112017 (2025).
5. S. N. Lowry *et al.*, “Spatial polarization modulation for terahertz single-pixel imaging,” *IEEE Trans. Terahertz Sci. Technol.* **14**, 386 (2024).
6. S. Zapata Valencia *et al.*, “Quantitative phase imaging in Hadamard-based active single-pixel microscopy by motionless transport of intensity equation,” *Opt. Lett.* **50**, 4478 (2025).
7. N. Radwell *et al.*, “Single-pixel infrared and visible microscope,” *Optica* **1**, 285 (2014).
8. A. Vallés *et al.*, “Broadband high-resolution terahertz single-pixel imaging,” *Opt. Express* **28**, 28868 (2020).
9. Z. Zhang *et al.*, “Hadamard single-pixel imaging versus Fourier single-pixel imaging,” *Opt. Express* **25**, 19619 (2017).
10. H. Tobón-Maya *et al.*, “Autofocusing method for active Hadamard single-pixel microscopy using gradient descent algorithms,” *Opt. Lasers Eng.* **185**, 108699 (2025).
11. J. Zhou and J. Yang, “Compressive sensing in image/video compression: sampling, coding, reconstruction, and codec optimization,” *Information* **15**, 75 (2024).
12. X. Yu *et al.*, “Deep compressive single pixel imaging by reordering Hadamard basis: a comparative study,” *IEEE Access* **8**, 55773 (2020).
13. M. F. Duarte *et al.*, “Single-pixel imaging via compressive sampling: building simpler, smaller, and less-expensive digital cameras,” *IEEE Signal Process Mag.* **25**, 83 (2008).
14. S.-X. Huang *et al.*, “Enhancing compressive single-pixel imaging with zig-zag-ordered Walsh-Hadamard light modulation,” *IEEE Photonics Technol. Lett.* **36**, 803 (2024).
15. X. Yu *et al.*, “Super sub-Nyquist single-pixel imaging by total variation ascending ordering of the Hadamard basis,” *Sci. Rep.* **10**, 9338 (2020).
16. Y. Cai *et al.*, “A detail-enhanced sampling strategy in Hadamard single-pixel imaging,” *Chin. Opt. Lett.* **21**, 071101 (2023).
17. A. L. Mur *et al.*, “Single-pixel image reconstruction from experimental data using neural networks,” *Opt. Express* **29**, 17097 (2021).
18. C. F. Higham *et al.*, “Deep learning for real-time single-pixel video,” *Sci. Rep.* **8**, 2369 (2018).
19. Z. Li *et al.*, “Single-pixel imaging with untrained convolutional autoencoder network,” *Opt. Laser Technol.* **167**, 109710 (2023).
20. F. Wang *et al.*, “Single-pixel imaging using physics enhanced deep learning,” *Photonics Res.* **10**, 104 (2022).
21. C. Li *et al.*, “An efficient augmented Lagrangian method with applications to total variation minimization,” *Comput. Optim. Appl.* **56**, 507–530 (2013).
22. S. Becker, J. Bobin, and E. J. Candès, “NESTA: a fast and accurate first-order method for sparse recovery,” *SIAM J. Imaging Sci.* **4**, 1 (2011).
23. S. van der Walt *et al.*, “scikit-image: image processing in Python,” *PeerJ* **2**, e453 (2014).
24. S. Węglarczyk, “Kernel density estimation and its application,” in *ITM Web of Conferences* (2018), Vol. **23**, 00037.
25. Y.-C. Chen, “A tutorial on kernel density estimation and recent advances,” *Biostat. Epidemiol.* **1**, 161 (2017).
26. A. Z. Zambon and R. Dias, “A review of Kernel density estimation with applications to econometrics,” *Int. Econom. Rev. (IER)* **5**, 20 (2013).

27. Y. Zheng and J. M. Phillips, " L_∞ error and bandwidth selection for Kernel density estimates of large data," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, 2015), p. 1533.
28. A. Gramacki, "*Selected applications related to Kernel density estimation*," (Springer, 2018), p. 133.
29. B. Ehteshami Bejnordi *et al.*, "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *JAMA* **318**, 2199 (2017).
30. B. S. Veeling *et al.*, "*Rotation equivariant CNNs for digital pathology*," (Springer, 2018), p. 210.
31. F. Pedregosa *et al.*, "Scikit-learn: machine learning in python," *J. Mach. Learn. Res.* **12**, 2825 (2013).
32. scikit-learn, "Density estimation," <https://scikit-learn.org/stable/modules/density.html>.
33. H. Tobón-Maya *et al.*, "Autofocusing method for active Hadamard single-pixel microscopy using gradient descent algorithms," *Opt. Lasers Eng.* **185**, 108699 (2025).