2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

# Spectral Security Imaging System (SSIS): Optical Authenticity for Hyperspectral Pushbroom Imaging

Pablo A Gomez T[1], Javier A Quiroga T[1], Hans Garcia[1], Gonzalo R. Arce[3], and Henry Arguello Fuentes[2]

*Abstract*— **Ensuring authenticity of hyperspectral imagery (HSI) at the moment of acquisition is critical: subtle spectral attacks can mislead downstream analysis before digital defenses take effect. By injecting the optical key before digitization, the effect is created in hardware and cannot be replicated in software, resulting in stronger protection. We present the Spectral Security Imaging System (SSIS), an acquisition-stage approach that injects a data-aware additive spectral key in the optical path of a pushbroom imager, binding integrity to the measurement while preserving the class-informative structure. We describe the complete system forward model, detector–key joint optimization, and a laboratory prototype together with a thorough calibration process. A laboratory dataset (unsigned and optically signed cubes) supports two evaluations. For manipulation detection, SSIS achieves detection accuracy of 92% with visual distortion of PSNR = 41.5 dB and SSIM = 0.981. For downstream classification on clean data, macro-F1 remains close to the unsigned ceiling, about 93% of the monochromator baseline (0.915 vs 0.981) and about 99% of the pushbroom baseline (0.892 vs 0.903), while outperforming multiplicative and watermarking baselines by up to 16.1 points in macro-F1 and 19.4 points in accuracy.**

## I. Introduction

Hyperspectral imaging (HSI) has become a cornerstone in remote sensing because of its ability to capture fine-grained spectral signatures across hundreds of contiguous wavelength bands. These spectral data are fundamental for critical applications such as environmental monitoring [1], [2], precision agriculture [3], climate studies, and disaster response [4]. However, the integrity of hyperspectral data is often threatened by unauthorized manipulations, which can compromise decision-making pipelines that depend on reliable spectral information. Examples include band-limited spectral shifts (e.g., red-edge/NIR offsets), spatial block edits that splice spectra across materials, cross-scene spectral splicing, band dropout or re-ordering, and radiometric gain/offset tampering in the digital pipeline. Even small spectral alterations can mislead classifiers, distort retrieval algorithms, and ultimately produce erroneous conclusions in downstream remote sensing tasks [5], [6]. Traditional approaches for ensuring hyperspectral data authenticity typically rely on post-acquisition mechanisms such as cryptographic hashing [7], [8], or software-based watermarking [9]–[12]. Although effective in detecting attacks in digital images, these approaches leave a critical vulnerability: they act after digitization, when the data may already have been modified. In remote sensing platforms, once the sensor output is digitized, adversaries can exploit digital channels to alter spectral content before security mechanisms are applied. As a result, post-acquisition strategies cannot fully guarantee the authenticity of the sensed spectra. In practice, cubes traverse firmware and digital pipelines (onboard processing, storage, downlink), where subtle band-dependent edits can cause silent downstream errors. Because post-acquisition protection only verifies an already-digitized cube, it cannot bind authenticity to analog acquisition effects (integration and ADC clipping/quantization), motivating pre-digitization key injection. Recent work has sought to complement data protection closer to the sensing stage. Optical watermarking and spectral keying methods introduce signatures directly in the light path of the imaging system, providing a degree of pre-digitization integrity [13]–[16]. For instance, acquisition-stage spectral key frameworks embed an optical key into a pushbroom sensor, enabling later detection of manipulated spectra [17]. Although promising, these approaches have focused primarily on algorithmic frameworks for detection, often overlooking the complete implementation and calibration of physical systems [8], [18].

In this work, we introduce the **Spectral Security Imaging System (SSIS)**, a spectral pushbroom-based imaging system that physically embeds spectral authenticity into the acquisition process. SSIS integrates a secondary device screen to project a spectral key into the optical path, ensuring that all acquired data carry a verifiable physical code. Crucially, the key is injected before digitization and is combined with the scene by the optics and sensor. Because this combination happens in hardware, it cannot be reproduced later by post-processing on a saved cube; a software watermark cannot recreate those physical statistics. SSIS emphasizes the end-to-end realization of an authenticated imaging system, including optical design, mathematical modeling, laboratory characterization, dataset acquisition, and task evaluation, as shown in Fig. 1. Unlike multiplicative signing, the additive key acts as a low-energy offset aligned with data statistics, which preserves class-discriminative structure under standard normalization. The result is a fully functional system that guarantees authenticity against post-processing manipulations while maintaining high spectral fidelity for remote sensing applications. Beyond imaging integrity, SSIS is also relevant to robotic perception pipelines, where hyperspectral data may be used for autonomous monitoring, mapping, or field inspection. In such settings, the detector can op-

[1]Dept. of E3T Eng., UIS, Bucaramanga, Colombia.
[2]Dept. of Systems Eng., UIS, Bucaramanga, Colombia.
[3]Dept. of ECE, Univ. of Delaware, Newark, DE, USA.
Corresponding author: Pablo Gomez. pablo2228330@correo.uis.edu.co

*This work was supported by the SENA and the MINISTERIO DE CIENCIA, TECNOLOGIA E INNOVACIÓN, under Convenio 268-2025.

erate as an acquisition-stage verification module that flags manipulated data before it is consumed by downstream classification, mapping, or decision-making blocks, thereby reducing the risk of silent perception failures. The main contributions of this work are:

- We propose the **Spectral Security Imaging System (SSIS)**, an optical pushbroom system that embeds a spectral signature using a codified device, ensuring authenticity at the point of acquisition.
- We develop a rigorous **forward model** of SSIS, acquiring the light propagation, dispersion, and signature embedding process, enabling both simulation and analytical evaluation of the system.
- We perform a complete **calibration and characterization** of the system, including wavelength-to-pixel mapping, spectral response, and optical throughput.
- We acquire a dedicated **laboratory hyperspectral dataset** signed with SSIS, providing a benchmark for authenticity research.
- We evaluate SSIS in two downstream tasks: (*i*) **manipulation detection** under several attacks, and (*ii*) **classification**, comparing signed and unsigned data to demonstrate robustness improvements.

## II. Related Work

Authentication for HSI divides into two lines: software methods that operate after digitization and optical approaches that integrate authenticity during acquisition. For clarity, $\mathbf{F} \in \mathbb{R}^{M\times N\times L}$ denotes the unsigned hyperspectral cube, and $\hat{\mathbf{G}}$ the digitized measurement used by prior art. The symbol $\mathbf{G}$ is reserved for our signed forward model in Sec. III.

### A. Post-Processing Authentication

*1) Encryption and Steganography:* Encryption treats the digitized hyperspectral cube $\hat{G}$ as a payload and hides it with an encryption operator $\mathscr{E}_k$ under a secret key $k$, producing an encrypted stream $\widetilde{G}$ for storage or transmission. This protects confidentiality and integrity in transit, but it occurs strictly *after* digitization: edits introduced before the sensor readout cannot be attested, and the distribution chain (compression, resampling, bit depth) may reshape inter-band relations that HSI analysis relies on, so encryption alone does not preserve measurement authenticity nor spectral semantics [19], [20]. Optical/digital image encryption for remote sensing and HSI is well studied [13], [15], [21], [22], but these methods still operate post-capture.

$$\widetilde{G} = \mathscr{E}_k(\hat{G}), \tag{1}$$

A related post-capture approach is steganography, which embeds a small message (e.g., a provenance hash) inside the image under a low-distortion budget. While useful for covert metadata, it likewise cannot attest pre-digitization manipulations, and making the tag robust typically requires stronger embedding that redistributes energy across bands and can bias spectral statistics used by classifiers [9], [14], [23], [24]. In contrast, our acquisition-stage additive optical signing binds integrity at capture and is designed to preserve class-informative spectra.

*2) Digital Watermarking:* Digital watermarking embeds an imperceptible signature in $\hat{G}$ to enable tamper evidence [7], [9], [10]. In the spatial domain, pixel intensities are perturbed under energy or PSNR budgets; in the transform domain, the embedding targets DCT, wavelet, or Fourier coefficients, while SVD-based schemes modify singular values to trade robustness against common edits with minimal visible impact [14], [25]–[27]. A compact abstraction is:

$$\hat{G}' = \mathscr{T}_u(\hat{G}) + \alpha\mathscr{U}, \tag{2}$$

where $\mathscr{T}_u$ is a transform (DCT, DWT, FT, SVD), $\mathscr{U}$ the embedded pattern, and $\alpha > 0$ its strength. Although effective for post-acquisition verification, stronger embeddings risk subtle spectral drifts that degrade classification, especially when the watermark is not aligned with class-discriminative spectral structure [14], [25]–[27].

### B. Acquisition-Based Optical Authentication

Optical approaches integrate the mechanism into the sensing chain so that integrity accompanies acquisition. In HSI, the prevailing formulation is multiplicative spectral modulation at acquisition. Using $F$ and a per-band key $c \in \mathbb{R}^L$ (or its spatially replicated tensor $C$), the digitized measurement is modeled as:

$$\hat{\mathbf{G}} = \mathbf{F} \odot \mathbf{C} + \mathbf{W}, \qquad C[m,n,\ell] \approx c_\ell, \tag{3}$$

with Hadamard (element-wise) product $\odot$ and measurement noise $W$. This pattern appears in wavelength-coded optics and spectral keying pipelines that target authenticity by altering band responses in a controlled fashion [13]–[16], [21], [22], [28]–[30]. A representative hyperspectral formulation adopts (3) with data-agnostic keys, initialized at random without adapting to the dataset's spectral energy profile [17]. While multiplicative modulation can increase detectability, it also changes the effective spectral transfer if $|c_\ell - 1|$ is not tightly constrained, which may reduce classification accuracy or require aggressive normalization at test time. Moreover, when the key is not aligned with data statistics, the balance between imperceptibility and verification power is not explicitly optimized for the target distribution, leading to sensitivity to calibration and suboptimal performance in both manipulation detection and downstream classification. Critically, the modulation must act in the continuous optical domain before sampling, because a pushbroom sensor integrates irradiance over exposure time, slit aperture, and spectral bandpass before A/D conversion, and the measurement is then shaped by sensor saturation and clipping, shot and read noise, analog gain, and quantization, all determined by the hardware rather than the stored cube. A post hoc digital modulation applied to a quantized datacube cannot reproduce these pre-ADC integrals and nonlinearities, therefore the authentication key must be realized physically in the acquisition path, not synthesized in post-processing. The above landscape motivates our approach in Sec. III: an additive, data-aware optical key designed so that detection leverages the dataset's spectral statistics while preserving class-informative structure for classification, supported by a calibrated forward model and laboratory characterization.

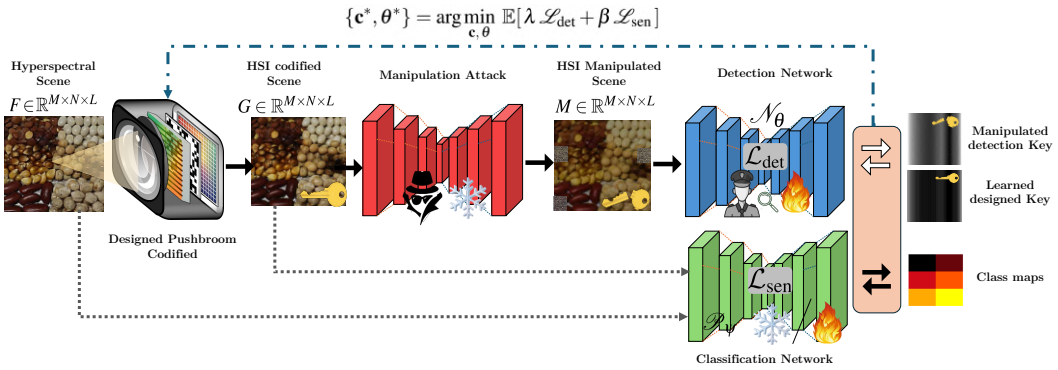


Fig. 1. SSIS pipeline and joint training. The unsigned cube $F$ is optically signed in the pushbroom via an additive spectral key $C = \Lambda c$ to yield $G$. Software attacks (block and band-limited spectral shift) generate $M$. The detector $\mathscr{N}_\theta$ regresses a 1-D key and drives the detection loss $\mathscr{L}_{\text{det}}$ (match on $G$, mismatch on $M$); a frozen classifier $\mathscr{P}_\psi$ provides the imperceptibility loss $\mathscr{L}_{\text{sen}}$ through the F1 gap between $F$ and $G$. We learn $c$ and $\theta$ by minimizing $\mathbb{E}[\lambda \mathscr{L}_{\text{det}} + \beta \mathscr{L}_{\text{sen}}]$.

## III. Proposed Optical Authenticity System

### A. Forward Model

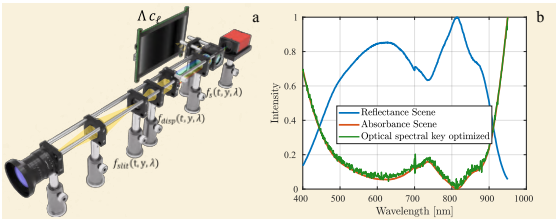


Fig. 2. SSIS optical authenticity system. (a) 3D layout of the pushbroom testbed implementation of the forward model with additive spectral–key injection. (b) Dataset spectra (scene reflectance/absorbance) together with the optimized optical key $k(\lambda) = \Lambda c(\lambda)$ used in our experiments.

We model SSIS using the same scene notation as above. Let $f(x,y,\lambda)$ denote the hyperspectral radiance of the scene, with $(x,y) \in \mathbb{R}^2$ and $\lambda \in \mathbb{R}_+$. A narrow slit selects a vertical strip while the platform scans along the $x$-axis at speed $v$, so that $x_t = vt$ and the slit output is $f_{\text{slit}}(t,y,\lambda) = f(x_t,y,\lambda)$. A dispersive element, prism, maps wavelength to a sensor coordinate along $y$ according to the dispersion law $\alpha(\lambda)$, yielding:

$$f_{\text{disp}}(t,y,\lambda) = f_{\text{slit}}\big(t, y - \alpha(\lambda), \lambda\big). \tag{4}$$

At this stage, SSIS introduces an additive spectral key in the optical path via a lateral electronic device channel coupled to an active-illumination screen, as shown in Fig. 2 a). To reflect the tunable illumination used in the prototype, we scale the key by a global intensity gain $\Lambda \in \mathbb{R}_+$ that is directly set in the implementation, while the spectral shape is given by a nonnegative function $c(\lambda)$. The signed optical field is therefore

$$f_s(t,y,\lambda) = f_{\text{disp}}(t,y,\lambda) + \Lambda c(\lambda), \tag{5}$$

where $c(\lambda)$ is held constant during each line exposure. The pushbroom readout integrates over the line period $\Delta t$, the pixel pitch $\Delta y$, and the effective bandpass $\Delta\lambda$ induced by the disperser and fore optics. Denoting by $m = 1,\dots,M$ the line index, $n = 1,\dots,N$ the across-track pixel index, and $\ell = 1,\dots,L$ the band index, the digitized, signed measurement $G \in \mathbb{R}^{M\times N\times L}$ is

$$G[m,n,\ell] = \int_{m\Delta t}^{(m+1)\Delta t} \int_{n\Delta y}^{(n+1)\Delta y} \int_{\ell\Delta\lambda}^{(\ell+1)\Delta\lambda} f_s(t,y,\lambda)\, \mathrm{d}\lambda\, \mathrm{d}y\, \mathrm{d}t. \tag{6}$$

Defining the band-averaged key $c_\ell = \int_{\ell\Delta\lambda}^{(\ell+1)\Delta\lambda} c(\lambda)\, \mathrm{d}\lambda$ and letting $F \in \mathbb{R}^{M\times N\times L}$ denote the unsigned cube induced by $f_{\text{disp}}$, we obtain the compact signed model

$$G \;=\; F + C + W, \qquad C[m,n,\ell] = \Lambda c_\ell, \tag{7}$$

i.e., the key is spectrally structured and replicated across space within each line exposure. Equivalently, using mode-1 Kronecker replication,

$$C \;=\; \Lambda J \otimes_1 (\mathbf{1}_N \otimes c), \tag{8}$$

with $J \in \mathbb{R}^{M\times 1\times 1}$ an all-ones tensor, $\mathbf{1}_N$ the $N$-vector of ones, and $c = [c_1,\dots,c_L]^\top$. The implementation assumes a line-synchronous key and a codified device settling time $\tau_{\text{rise}}$ smaller than the exposure $T_{\text{exp}}$ to avoid intra-line transients. The detector responsivity $r(\lambda)$ over the bands used is stable during acquisition; we pre-normalize measurements so that $r(\lambda)$ is absorbed into $F$ and $c(\lambda)$, and any residual drift is captured by $W$. Finally, $\Lambda$ is chosen to keep spectral distortion small, which preserves class-informative structure for downstream classification.

### B. Hyperspectral Manipulations

We evaluate two families of attacks simulated in software on the signed SSIS laboratory cubes $G$ and let $M \in \mathbb{R}^{M\times N\times L}$ be the manipulated cube. Class masks $\mathbf{O}_{cls}[m,n]$ come from the SSIS laboratory dataset. We focus on spatially coherent edits and band-limited spectral shifts.

**(1) Spatially coherent block attack ($M_{blk}$).** We alter only $s \times s$ patches inside a target class which are classes from the same image to enforce spatial coherence. Define an edit mask $E[m,n] \in \{0,1\}$ that flags the selected blocks within $O_{cls}$. For each edited pixel we assign another reference spectrum $S[m,n,\ell]$ taken from the closest pixel that belongs to a different class $O_{cls}$ in the same scene. With a blend factor $\rho \in [0,1]$,

$$M_{\text{blk}}[m,n,\ell] = \begin{cases} G[m,n,\ell] & \text{if } E[m,n] = 0, \\ \rho\, G[m,n,\ell] + (1-\rho)\, S[m,n,\ell] & \text{else.} \end{cases} \tag{9}$$

This produces contiguous regions whose broadband look remains plausible while spectra are replaced by class-inconsistent signatures.

**(2) Band-limited spectral shift ($M_{spec}$).** We perturb a contiguous set of bands for pixels of a target class while keeping all other bands unchanged. Let $\mathscr{B} \subset \{1,\dots,L\}$ be a band window and let $\delta$ be a small offset:

$$M_{\text{spec}}[m,n,\ell] = \begin{cases} G[m,n,\ell] + \delta, & \text{if } O_{\text{cls}}[m,n] = 1 \ \wedge\ \ell \in \mathscr{B}, \\ G[m,n,\ell], & \text{else.} \end{cases} \tag{10}$$

Consistent with our instrument's coverage, we consider three windows widely used in agriculture: base classify (520–600 nm), Red-Edge (700–740 nm), and NIR up to the sensor limit (800–850 nm).

### C. Optical Spectral Key Optimization

The spectral key $c \in \mathbb{R}^L_+$ scales additively into the sensing path via a global intensity gain $\Lambda$ fixed by calibration, so that $C[m,n,\ell] = \Lambda c_\ell$. The key is optimized to improve attack detection while remaining visually and spectrally unobtrusive. Training uses unsigned cubes $F$, their signed versions, and software-manipulated cubes $M$ generated with the attacks in Sec. III-B. An auxiliary optical modulation path provides the physical channel where the key is injected. The key is line-synchronous and constant within each exposure, matching the forward model in Sec. III-A.

*a) Detection network and loss:* The detector $\mathscr{N}_\theta(\cdot) : \mathbb{R}^{M\times N\times L} \to \mathbb{R}^L$ estimates a one-dimensional key in band space. Let the target vector be $\mathbf{k} = \Lambda c$. The detection loss encourages two complementary behaviours: (i) on clean signed data $G$, the network should recover the true key as closely as possible; and (ii) on manipulated data $M$, the network's estimate should deviate from the true key, signalling that the content has been altered. $\mathscr{N}_\theta$ aggregates the cube spatially (global average pooling over $m,n$) and outputs a band vector in $\mathbb{R}^L$. Concretely,

$$\mathscr{L}_{\text{det}} = \frac{\|\mathbf{k} - \mathscr{N}_\theta(G)\|_2}{\|\mathbf{k}\|_2} + \left(1 - \frac{\|\mathbf{k} - \mathscr{N}_\theta(M)\|_2}{\|\mathbf{k}\|_2 + \varepsilon_{\text{n}}}\right)^2. \tag{11}$$

This loss focuses and sharpens the margin between clean and attacked inputs and improves detection precision.

*b) Imperceptibility loss (F1-gap):* To link imperceptibility to task performance, we penalize the gap between the classification F1 obtained on the unsigned cube and on the signed cube. Using a fixed and pretrained classifier $\mathscr{P}_{\psi^*}$ and ground-truth labels $O_{\text{cls}}$ over the supervised pixel set $\Omega$, let $\text{F1}(\cdot)$ denote the macro-F1 score. The imperceptibility loss:

$$\mathscr{L}_{\text{sen}} = \Big(\text{F1}\big(\mathscr{P}_{\psi^*}(F), O_{\text{cls}}; \Omega\big) - \text{F1}\big(\mathscr{P}_{\psi^*}(G), O_{\text{cls}}; \Omega\big)\Big)^2. \tag{12}$$

Minimizing (12) drives the key to preserve the classifier's effectiveness, keeping the signed-data F1 close to the unsigned baseline.

**Algorithm 1** Key Optimization with Detection (compact, $k = \Lambda c$ in losses)

**Require:** Unsigned cubes $F$ with labels $O_{\text{cls}}$ on $\Omega$; epochs $T$; inner iters $K$; step size $\alpha$; weights $\lambda, \beta$; fixed $\Lambda$; attack generator SimAttacks$(\cdot)$; frozen classifier $\mathscr{P}_\psi$

Initialize key $c^0$ (data-aware); project to $\|c^0\|_\infty \le \varepsilon$, $c^0 \ge 0$; initialize detector $\theta^0$
**for** $t = 1$ to $T$ **do**
  Sample mini-batch $(F, O_{\text{cls}})$
  Build $C$ with $C[m,n,\ell] = \Lambda c_\ell^{t-1}$; set $G \leftarrow F + C$
  Set $k \leftarrow \Lambda c^{t-1}$ ▷ $k \in \mathbb{R}^L$
  Generate $M \leftarrow \text{SimAttacks}(G)$
  **for** $k_{\text{in}} = 1$ to $K$ **do**
    $\mathscr{L}_{\text{det}} \leftarrow \frac{\|k - \mathscr{N}_\theta(G)\|_2}{\|k\|_2} + \Big(1 - \frac{\|k - \mathscr{N}_\theta(M)\|_2}{\|k\|_2}\Big)^2$
    $\mathscr{L}_{\text{sen}} \leftarrow \Big(\text{F1}\big(\mathscr{P}_\psi(F), O_{\text{cls}}; \Omega\big) - \text{F1}\big(\mathscr{P}_\psi(G), O_{\text{cls}}; \Omega\big)\Big)^2$
    $\mathscr{L} \leftarrow \lambda\, \mathscr{L}_{\text{det}} + \beta\, \mathscr{L}_{\text{sen}}$
    $\theta^t \leftarrow \theta^{t-1} - \alpha \nabla_\theta \mathscr{L}$
    $\tilde{c} \leftarrow c^{t-1} - \alpha \nabla_c \mathscr{L}$
    $C[m,n,\ell] \leftarrow \Lambda c_\ell^t$ ; update $C$ accordingly
Output: $c^*$, $\theta^*$

*c) Overall objective and optimization:* The detector $\mathscr{N}_\theta$ and the spectral key $c$ are trained jointly, while the classifier $\mathscr{P}_\psi$ remains fixed. The classifier is pretrained on unsigned data and is not updated during key optimization. The total loss combines the detection loss $\mathscr{L}_{\text{det}}$ and the imperceptibility loss $\mathscr{L}_{\text{sen}}$, where $\mathscr{L}_{\text{sen}}$ is the F1-gap between unsigned and signed inputs computed with $\mathscr{P}_\psi$:

$$\{\mathbf{c}^*, \theta^*\} = \arg\min_{\mathbf{c},\theta} \mathbb{E}[\,\lambda\, \mathscr{L}_{\text{det}} + \beta\, \mathscr{L}_{\text{sen}}\,], \tag{13}$$

subject to $c \ge 0$, and fixed $\Lambda$. With hyperparameters $\lambda, \beta > 0$. After each gradient step, $c$ is projected onto a feasible set (nonnegativity and a small $\ell_\infty$ budget by clipping methodology) to satisfy imperceptibility and hardware limits, as shown in Alg. 1. Attacks are simulated in software during training to generate manipulated inputs for the detection loss.

## IV. Optical System Implementation

Building on the forward model and the key optimization described in Sec. III, we build a laboratory pushbroom prototype that implements the additive signing term $C[m,n,\ell] = \Lambda c_\ell$ directly in the optical path, as shown in Fig 3. A static scene is imaged onto a vertical slit, relayed through the dispersive stage, and recorded on the detector while the scan motion provides the along-track sampling $\Delta x = v\Delta t$. A programmable optical encoder is combined through a beam path to inject a line-synchronous spectral key that is uniform across the slit during each exposure and varies only with wavelength according to the vector $c$. The global gain $\Lambda$ is set once during calibration and kept fixed for all acquisitions, whereas $c$ is selected by the optimization procedure and held constant within each exposure.

The detector is operated in its linear regime, and raw frames are stacked over time to form the hyperspectral cube. Prior to the calibration procedures, we center the slit on the sensor, align the dispersion axis to the pixel grid, focus across the wavelength range by minimizing line widths at a few references, and verify that the injected signal is

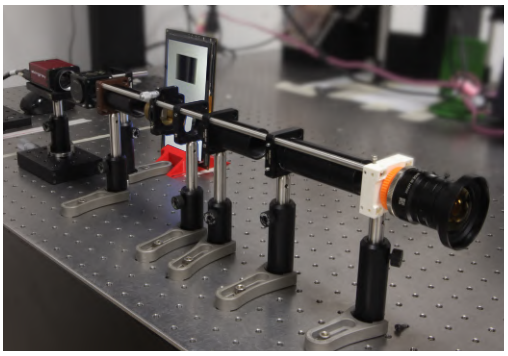

Fig. 3. Laboratory pushbroom testbed implementing the additive spectral-key injection.

spatially uniform across the slit. These steps establish the operating point used in the next section, where we calibrate wavelength, spatial scale, and radiometry.

## V. Calibration and Characterization

This section reports the calibration procedures needed to operate the system in a repeatable and analyzable regime. We first map detector coordinates to wavelength and quantify spectral dispersion and bandpass, including smile and keystone. We then establish spatial scale and imaging performance in both scan and dispersion directions, followed by the spatial registration between the encoder plane and the sensor to guarantee that the per-band key is applied consistently. Next, we set a radiometric operating point by calibrating illumination and exposure, estimating noise and SNR per band, and tying the global gain $\Lambda$ to physical counts. All notation follows the definitions introduced in Sec. III.

### A. Wavelength Calibration and Spectral Dispersion

Figure 4 summarizes the spectral calibration of our pushbroom prototype and provides the quantities needed to register wavelength to detector coordinates and to estimate band separability. Panel (a) reports the wavelength–to–pixel mapping obtained from the TLS-300XR sweep; the mapping is monotonic and its slope decreases toward the NIR (e.g., $\sim$ 0.36 px/nm over 401–463 nm versus $\sim$ 0.09 px/nm over 632–849 nm, an $\sim 4\times$ reduction in dispersion per nanometer). The low-order fit overlaid in (a) is used throughout to map wavelengths to detector coordinates, while panel (c) normalizes the measured line profiles to visually emphasize how overlap increases as dispersion diminishes toward longer wavelengths. Panel (b) characterizes the instrument bandpass via the line profiles and their FWHM, and compares it to the spacing between adjacent centroids to assess separability. In pixels, the FWHM increases from $\sim 2.0$ px at 401 nm to $\sim 6.9$ px at 848.6 nm ($\approx 3.4\times$); expressed in nanometers, the broadening is larger (from $\sim 2.32$ nm to $\sim 50.33$ nm) due to the reduced px/nm conversion at longer wavelengths. Using a conservative Rayleigh-like non-overlap criterion, we estimate 51 separable bands over 400–850 nm, with the highest band density between 400 and 500 nm, consistent with the overlap trends visualized in (c).

### B. Spatial Calibration and Imaging Performance

We characterized the imaging point spread function (PSF) across field and wavelength using the TLS-300XR as a narrowband source. A diffraction-limited spot was formed at the object plane and imaged through the full pushbroom train under the same focus and aperture used for data acquisition. We sampled three wavelengths within our band, $\{550, 650, 750\}$ nm, and for each one we acquired patches at three lateral field locations. Each patch was recentred with subpixel registration and normalized to the per-row peak to isolate shape changes from illumination drift. Figure 5 shows the resulting mosaic, with a 50 px scale bar as a reference.

From Fig. 5 we observe a stable PSF over the field: the central and lateral patches have similar footprint and symmetry, with only mild edge broadening and haloing at the left/right columns. This indicates a weak field dependence, consistent with small residual astigmatism or field curvature, but not large enough to impact the patch sizes used in our experiments. Spectrally, the PSF shows a tightening as wavelength increases, while maintaining its elongated shape along the slit image; this trend matches the expectation of slightly improved MTF at longer wavelengths for our sensor-optic combination. Within the calibrated operating point the PSF is effectively space–invariant across the central field and only weakly wavelength–dependent, which supports the single–kernel assumption in the forward model.

## VI. Laboratory Scene and Spectral Signatures

To target an agricultural use case, we acquired a representative laboratory scene under the calibrated operating point (Sec. IV). For this scene we captured an unsigned cube $F$ and its signed counterpart given by Eq. 7, using identical illumination, exposure, and $\Lambda$. The scene contains six seed/food classes with mixed textures and annotated regions for evaluation. Figure 6 summarizes the data. The false-color composite shows uniform coverage, and the class-wise mean spectra ($\pm 1\sigma$) exhibit clear separability with consistent radiometric trends: corn/peas present higher intensity in the 520–600 nm range, while beans/lentils are dimmer and flatter, and chickpeas show a distinct rise toward the red edge. Overlaying signed and unsigned means yields near-coincident curves across the band, indicating that the additive key preserves relative intensities and discriminative shapes while introducing only a small, low-energy bias.

## VII. Results

We assess SSIS in two complementary settings that connect the additive model in Eq. 7 with its physical realization: a monochromator sweep with Signed (sim.), which emulates ideal bandwise signing to validate the formulation, and a pushbroom sweep with Signed Imp. (ours), which implements the calibrated hardware in Sec. IV. Each setting is compared against its Unsigned reference and three baselines Spatial WM [11], DCT WM [31], and the optical multiplicative key Opt. Mult. [17]; for pushbroom we also contrast to the monochromator unsigned to gauge alignment with the ideal case. Manipulations follow Sec. III-B (Eqs. 9,

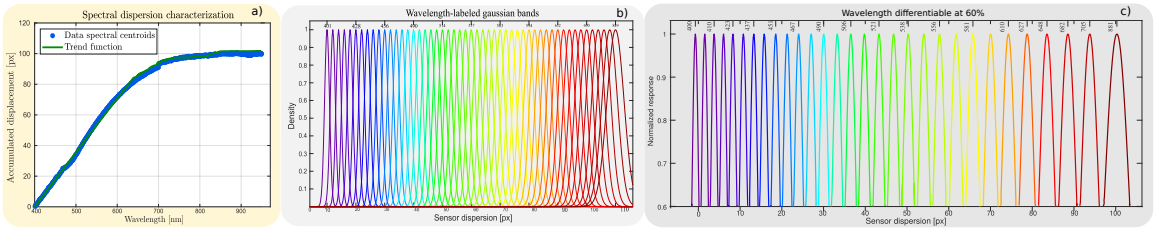


Fig. 4. Spectral characterization. **(a)** Wavelength–to–pixel mapping from the TLS-300XR sweep: measured centroids and the low-order fit used throughout the paper; dispersion per nanometer decreases toward the NIR. **(b)** Gaussian line profiles that model the instrument bandpass at each wavelength. Example full widths at half maximum (FWHM): at 410 nm $\approx$ 2.3 nm, at 521 nm $\approx$ 4.6 nm, and at 881 nm $\approx$ 20 nm. **(c)** Same profiles normalized; we adopt the 50% bandwidth (FWHM) as a non-overlap criterion to avoid adjacent-band interference, yielding 51 resolvable bands over 410–881 nm.

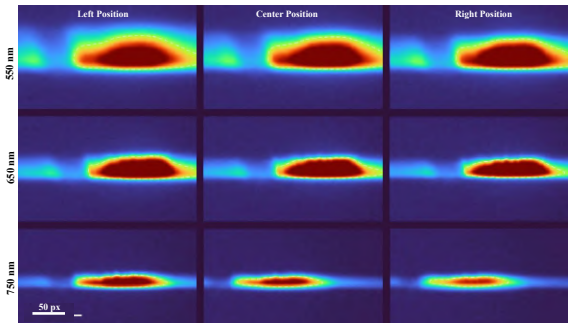


Fig. 5. PSF across field (columns: left/center/right) and wavelength (rows: 550, 650, 750 nm). Each patch is subpixel-centered and normalized; the 50 px bar provides a common spatial scale.

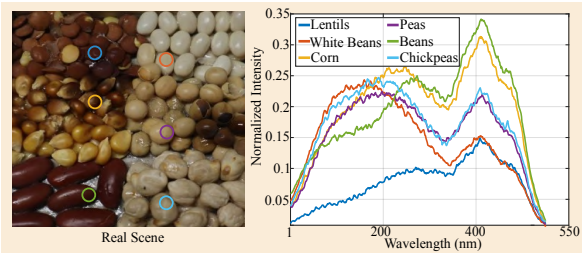


Fig. 6. Laboratory scene and spectra. Left: false-color composite of the unsigned cube $F$. Right: class mean spectral signatures with $\pm 1\sigma$ bands. Unsigned means (solid) and signed means from Eq. 7 (dashed) largely overlap, showing that signing preserves class-discriminative radiometry.

10); evaluation covers pixel-wise detection accuracy (ACC), visual fidelity (PSNR/SSIM) together with a probe-pixel SAM, and downstream classification (ACC, macro-F1) using a frozen classifier trained on Unsigned.

### A. Implementation and reproducibility details

SSIS is implemented in PyTorch and trained on a single NVIDIA RTX 4090 GPU. The detector $N_\theta$ maps an input cube in $\mathbb{R}^{M\times N\times L}$ to a bandwise key estimate in $\mathbb{R}^L$ using a lightweight 3D convolutional feature extractor, global average pooling, and a regression head. During training, cubes are formed from $F$, $C$, and $\zeta$, while manipulated samples $M$ are generated using the attacks in Sec. III-B. The detector parameters $\theta$ and the key $c$ are jointly optimized with Adam by minimizing $\lambda L_{\text{det}} + \beta L_{\text{sen}}$ for 1000 outer epochs and 500 inner updates per epoch. The optical prototype uses an Allied Vision Stingray F-080B monochrome camera in a slit-based pushbroom configuration, an Amici prism for spectral dispersion, and a broadband Navitar objective. Frames are acquired at a fixed calibrated operating point with an integration time of 1 $\mu$s per shot and stacked.

### B. Manipulation Attack Analysis

A parameter sweep was conducted for the two manipulation models in Sec. III-B (Eqs. 9, 10) under both acquisition settings: the monochromator sweep with Signed (sim.) and the pushbroom sweep with Signed Imp. (ours) (hardware in Sec. IV). For block modification, the patch size was varied as $s \in \{16, 32, 64\}$; for band-limited spectral shift, Green (520–600 nm), Red-Edge (700–740 nm), and NIR (800–850 nm) were evaluated with $\delta = 0.03$. Each condition is compared against Spatial WM, DCT WM, and the optical multiplicative key Opt. Mult. [17]. Pixel-wise detection accuracy (ACC) is reported with detector thresholds fixed during evaluation; joint key–detector training follows Sec. III-C. Results are summarized in Fig. 7.

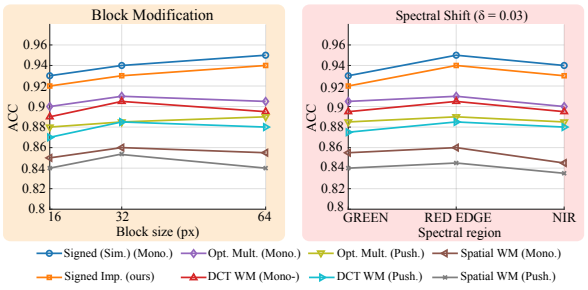


Fig. 7. Detection accuracy (ACC) under block modification (left) and spectral shift (right) for the monochromator sweep (Mono.) and the pushbroom implementation (Push).

From Fig. 7, ACC under block modification increases monotonically with patch size (typically by $\sim$2–3 pp from $s$=16 to $s$=64) for all methods. Across $s$, Signed (sim.) yields the top curve, and Signed Imp. (ours) tracks it closely; their Mono.–Push. gap remains within $\sim$0.3–1.5 pp. The additive signing maintains a consistent margin of $\sim$2–4 pp over Opt. Mult. and $\sim$6–10 pp over DCT/Spatial WM. Under spectral shift, the additive methods also lead in Green/Red-Edge/NIR; gaps versus Opt. Mult. are typically $\sim$3–4 pp in Red-Edge and $\sim$2–3 pp in Green/NIR, while margins over watermarking remain around $\sim$6–9 pp. ACC in NIR is lower by about 1–2 pp for all methods relative to Green/Red-Edge, yet the relative ordering is preserved. Overall, the

TABLE I

PERFORMANCE UNDER IDENTICAL ATTACKS FOR UNSIGNED AND SIGNED DATA. DETECTION IS REPORTED ONLY FOR SIGNED DATA, SINCE UNSIGNED DATA PROVIDES NO INTEGRITY SIGNAL.

| Attack | Param. | Det. (Signed) ↑ | F1 Uns. ↑ | F1 Sig. ↑ | SAM Uns. ↓ | SAM Sig. ↓ |
|---|---|---|---|---|---|---|
| Block | 16 | 0.920 | **0.870** | 0.861 | **0.034** | 0.037 |
| Block | 32 | 0.930 | **0.842** | 0.833 | **0.041** | 0.044 |
| Block | 64 | 0.940 | **0.811** | 0.801 | **0.052** | 0.056 |
| Shift | Green | 0.920 | **0.878** | 0.869 | **0.031** | 0.034 |
| Shift | Red Edge | 0.940 | **0.826** | 0.816 | **0.048** | 0.052 |
| Shift | NIR | 0.930 | **0.836** | 0.826 | **0.044** | 0.048 |

TABLE II

PIXEL-WISE CLASSIFICATION ON CLEAN TEST DATA. CLASSIFIER FROZEN (TRAINED ON UNSIGNED). MEAN±STD ACROSS SCENES.

| *Monochromator sweep* | | | *Pushbroom sweep (our implementation)* | | |
|---|---|---|---|---|---|
| Method | ACC ↑ | Macro-F1 ↑ | Method | ACC ↑ | Macro-F1 ↑ |
| **Unsigned** | **0.961±0.03** | **0.981±0.01** | Unsigned | 0.942±0.05 | 0.903±0.02 |
| **Signed (sim.)** | 0.929±0.02 | 0.915±0.02 | **Signed Imp. (ours)** | 0.911±0.02 | 0.892±0.02 |
| Spatial WM | 0.735±0.06 | 0.754±0.06 | Spatial WM | 0.718±0.05 | 0.738±0.05 |
| DCT WM | 0.809±0.02 | 0.805±0.02 | DCT WM | 0.802±0.03 | 0.818±0.03 |
| Opt. Mult. [17] | 0.830±0.04 | 0.840±0.03 | Opt. Mult. [17] | 0.822±0.03 | 0.842±0.03 |

pushbroom implementation reproduces the behavior of the ideal emulation with small Mono.–Push. differences, while consistently outperforming multiplicative signing and both watermarking baselines across attack regimes. We next assess visual fidelity and spectral consistency in Fig.8.

Table I shows that identical attacks degrade both unsigned and signed data, but only signed data provides an integrity signal through key verification. Thus, SSIS does not eliminate task degradation under attack; rather, it makes manipulation detectable and avoids silent downstream failures.

### C. Visual fidelity and spectral consistency

Under the calibrated operating point (Sec. IV) and the training setup of Sec. III-C, Fig. 8 compares composites for both acquisition settings: monochromator sweep (top row) and pushbroom sweep (bottom row). Columns correspond to Unsigned, Signed (sim.) [proposed, ideal emulation], Spatial WM, DCT WM, and Opt. Mult. [17]; the proposed columns are highlighted in orange. Each panel reports PSNR/SSIM with respect to the corresponding unsigned reference. A red marker indicates the probe pixel used to compute a SAM per-pixel; the right-hand plots display the probe spectra and the SAM values reported in the legend. From Fig. 8, the proposed additive signing yields higher image-space fidelity than the baselines in both settings, improving PSNR by about +2–6 dB and SSIM by about +0.004–0.038 over Opt. Mult. and watermarking at comparable distortion budgets. The pushbroom implementation preserves appearance closely relative to the unsigned ideal (monochromator) reference and tracks the Signed (sim.) column, indicating that the hardware reproduces the intended behavior of the emulated design. Spectrally, the key insertion does not alter the signature shape: the probe-pixel SAM remains small for the proposed methods $\lesssim$0.02 rad for Signed (sim.) and $\lesssim$0.03 rad for Signed Imp. (ours), with a minimum advantage of $\geq$0.02 rad versus the next-best baseline consistent with an additive bias that preserves class-informative structure.

### D. Classification Performance

Downstream utility on clean data is evaluated with the same baselines used above. The classifier is frozen and pretrained on unsigned cubes; at test time, only the acquisition/embedding differs across methods (Table II). Inputs are per-pixel spectra with a fixed spatial context window, per-band normalization, and cross-entropy training on the unsigned. We report overall accuracy (ACC) and macro-F1 over the annotated pixels $\Omega$.

From Table II, in the monochromator sweep the proposed Signed (sim.) improves over Opt. Mult. by +9.9 pp in ACC (0.929 vs. 0.830) and +7.5 pp in macro-F1 (0.915 vs. 0.840); the margins relative to DCT WM are +12.0 pp and +11.0 pp, and relative to Spatial WM are +19.4 pp and +16.1 pp. The distance to the unsigned reference is 3.2 pp in ACC (0.961 vs. 0.929) and 6.6 pp in macro-F1 (0.981 vs. 0.915). In the pushbroom sweep, the implemented Signed Imp. (ours) exceeds Opt. Mult. by +8.9 pp in ACC (0.911 vs. 0.822) and +5.0 pp in macro-F1 (0.892 vs. 0.842); relative to DCT WM the gains are +10.9 pp and +7.4 pp, and relative to Spatial WM they are +19.3 pp and +15.4 pp. The distance to its unsigned pushbroom reference remains small: 3.1 pp in ACC (0.942 vs. 0.911) and 1.1 pp in macro-F1 (0.903 vs. 0.892), indicating that the additive optical signing consistently outperforms multiplicative signing and watermarking while staying close to the unsigned ceiling in both settings.

## VIII. CONCLUSIONS

We presented SSIS, a pushbroom hyperspectral acquisition pipeline that embeds a data-aware additive spectral key in the optical path and binds integrity to the measurement at capture. Across both idealized and hardware acquisition settings, SSIS achieves the highest manipulation-detection accuracy while preserving spectral fidelity and downstream utility, outperforming optical multiplicative signing and watermarking baselines. These results show that optical signing can provide a practical end-to-end path to hyperspectral authenticity without compromising standard analysis pipelines. Our current evaluation is limited to a controlled laboratory operating point, and extending SSIS to in-the-wild data collection remains an important next step. Future work will broaden the evaluation to more diverse operating conditions, including changes in illumination, sensor gain/exposure, scene content, and temporal drift, as well as additional threat models such as band dropout/re-ordering, compression/resampling artifacts, cross-scene splicing, and adaptive adversarial attacks. In future work, a quantitative characterization of detector latency, memory footprint, and verification throughput will also be included to assess deployment readiness. We also plan to study key rotation and weak line-wise randomization to mitigate replay and key-estimation attacks, and to move the detector closer to the sensor for low-latency verification.

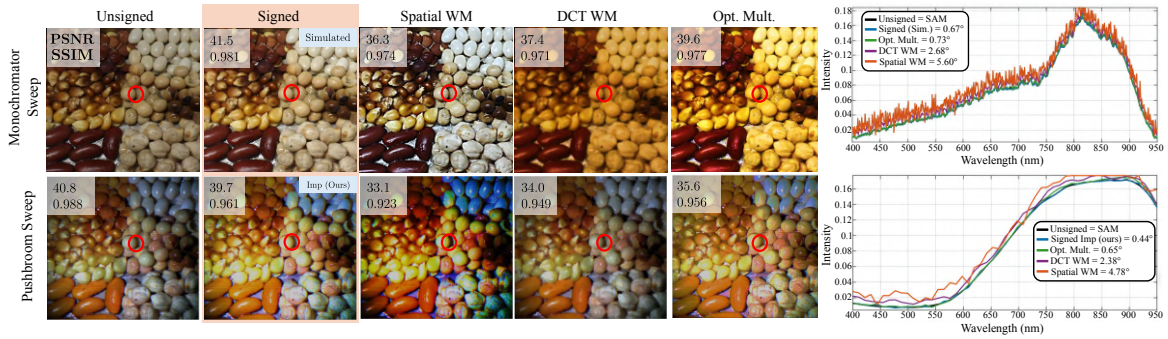


Fig. 8. False-color comparison with PSNR/SSIM relative to the unsigned reference and probe-pixel SAM. Top: monochromator sweep. Bottom: pushbroom sweep. Proposed methods are highlighted in orange: Signed (ideal emulation, top) and Signed Imp. (hardware implementation, bottom). The red circle marks the probe pixel used for the SAM trace. At comparable distortion budgets, the proposed additive signing preserves appearance and spectra

## REFERENCES


[1] A. Bhargava, A. Sachdeva, K. Sharma, M. H. Alsharif, P. Uthansakul, and M. Uthansakul, "Hyperspectral imaging and its applications: A review," *Heliyon*, vol. 10, no. 12, p. e33208, 2024.

[2] Q. Zhe, W. Gao, and e. a. Zhang, "A hyperspectral classification method based on deep learning and dimension reduction for ground environmental monitoring," *IEEE Access*, 2025.

[3] B. Lu, Dao, and et al, "Recent advances of hyperspectral imaging technology and applications in agriculture," *Remote Sensing*, vol. 12, no. 16, 2020.

[4] S. Faisal, M. P.-L. Ooi, S. K. Abeysekera, Y.-C. Kuang, and D. Fletcher, "Roadmap for measurement and applications: Uncertainty quantification and visualization for optimal decision-making in hyperspectral imaging-based precision agriculture," *IEEE Instrumentation & Measurement Magazine*, vol. 28, no. 1, pp. 23–32, 2025.

[5] M. Shimoni and e. a. Haelterman, "Hypersectral imaging for military and security applications: Combining myriad processing and sensing techniques," *IEEE Geoscience and Remote Sensing Magazine*, vol. 7, no. 2, pp. 101–117, 2019.

[6] K. Dasari, S. A. Yadav, L. Kansal, J. Adilakshmi, G. Kaliyaperumal, and A. Albawi, "Fusion of hyperspectral imaging and convolutional neural networks for early detection of crop diseases in precision agriculture," in *2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE)*. IEEE, 2024.

[7] P. Singh and R. S. Chadha, "A survey of digital watermarking techniques, applications and attacks," *International Journal of Engineering and Innovative Technology (IJEIT)*, vol. 2, no. 9, pp. 165–175, 2013.

[8] X. Li, M. Ding, Y. Gu, and A. Pižurica, "An end-to-end framework for joint denoising and classification of hyperspectral images," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 7, pp. 3269–3283, 2023.

[9] K. M. Hosny, A. Magdi, O. ElKomy, and H. Hamza, "Digital image watermarking using deep learning: A survey," *Computer Science Review*, vol. 53, p. 100662, 2024.

[10] A. Dubey, N. Dixit, and e. a. Arora, "Challenges and opportunities in digital image watermarking," in *2024 First International Conference on Pioneering Developments in Computer Science and Digital Technologies (IC2SDT)*, 2024, pp. 29–34.

[11] A. K. Singh, N. Sharma, M. Dave, and A. Mohan, "A novel technique for digital image watermarking in spatial domain," in *2012 2nd IEEE International Conference on Parallel, Distributed and Grid Computing*. IEEE, 2012, pp. 497–501.

[12] E. Yavuz and Z. Telatar, "Improved svd-dwt based digital image watermarking against watermark ambiguity," in *Proceedings of the 2007 ACM symposium on Applied computing*, 2007, pp. 1051–1055.

[13] H. Chen, C. Tanougast, and et al, "Optical hyperspectral image encryption based on improved chirikov mapping and gyrator transform," *Optics and Lasers in Engineering*, vol. 107, pp. 62–70, 2018.

[14] G. Qu, X. Meng, X. Yang, H. Wu, P. Wang, W. He, and H. Chen, "Optical color watermarking based on single-pixel imaging and singular value decomposition in invariant wavelet domain," *Optics and Lasers in Engineering*, vol. 137, p. 106376, 2021.

[15] H. Li, X. Bai, M. Shan, and et al, "Optical encryption of hyperspectral images using improved binary tree structure and phase-truncated discrete multiple-parameter fractional fourier transform," *Journal of Optics*, vol. 22, no. 5, p. 055701, apr 2020.

[16] M. Shan, J. Guo, and et al, "Improved multiple-image authentication based on optical interference by wavelength multiplexing," *Appl. Opt.*, vol. 61, no. 23, pp. 6931–6938, Aug 2022.

[17] P. Gomez, R. Jacome, E. Martinez, H. Garcia, and H. Arguello, "Optical authenticity in pushbroom system for spectral information protection," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.

[18] Q. Li, J. Li, T. Li, and Y. Feng, "A joint framework for underwater hyperspectral image restoration and target detection with conditional diffusion model," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 17 263–17 277, 2024.

[19] G. L. Friedman, "The trustworthy digital camera: Restoring credibility to the photographic image," *IEEE Transactions on consumer electronics*, vol. 39, no. 4, pp. 905–910, 1993.

[20] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.

[21] Y. Xiong, J. Du, and C. Quan, "Optical encryption and authentication scheme based on phase-shifting interferometry in a joint transform correlator," *Optics and Laser Technology*, vol. 126, p. 106108, 2020.

[22] H. Wei and X. Wang, "Optical multiple-image authentication and encryption based on phase retrieval and interference with sparsity constraints," *Optics and Laser Technology*, vol. 142, p. 107257, 2021.

[23] G. Fan, Z. Pan, Q. Zhou, J. Dong, and X. Zhang, "Reversible data hiding in multispectral images for satellite communications," *Journal of Information Security and Applications*, vol. 67, p. 103180, 2022.

[24] M. Bodke and S. Chaudhari, "Hyperspectral remote sensing image watermarking using discrete wavelet transform and forensic based investigation archimedes optimization," *Earth Science Informatics*, vol. 17, no. 5, pp. 4297–4313, Jul 2024.

[25] X. Wu, J. Hu, Z. Gu, and J. Huang, "A secure semi-fragile watermarking for image authentication based on integer wavelet transform with parameters," in *Conferences in Research and Practice in Information Technology Series*, vol. 108, 2005, pp. 75–80.

[26] E. Najafi, "A robust embedding and blind extraction of image watermarking based on discrete wavelet transform," *Mathematical Sciences*, vol. 11, pp. 307–318, 2017.

[27] M. Cedillo-Hernandez, F. Garcia-Ugalde, M. Nakano-Miyatake, and H. Perez-Meana, "Robust watermarking method in dft domain for effective management of medical imaging," *Signal, Image and Video Processing*, vol. 9, pp. 1163–1178, 2015.

[28] Y. Zhao, B. Liu, T. Zhu, M. Ding, X. Yu, and W. Zhou, "Proactive image manipulation detection via deep semi-fragile watermark," *Neurocomputing*, vol. 585, p. 127593, 2024.

[29] V. Asnani, X. Yin, T. Hassner, S. Liu, and X. Liu, "Proactive image manipulation detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 15 386–15 395.

[30] V. Asnani and e. a. Yin, "Malp: Manipulation localization using a proactive scheme," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 12 343–12 352.

[31] M. I. Khan, Rahman *et al.*, "Digital watermarking for image authentication based on combined dct, dwt and svd transformation," *arXiv preprint arXiv:1307.6328*, 2013.