Spectral Band Attribution in Historical Ink and Substrate Recognition

Marco Buzzelli [1], Ana Belén López-Baldomero [2], Francisco Moronta-Montero [2], Eva Maria Valero [2]
[1] Department of Informatics Systems and Communication, University of Milano-Bicocca,
Viale Sarca, 336, Milan, 20126, Italy
[2] Department of Optics, University of Granada, Faculty of Sciences,
Campus Fuentenueva, Granada, 18071, Spain

ABSTRACT

We present an interpretative approach to classifying inks and substrates in historical handwritten documents using hyperspectral imaging and deep learning. We introduce a custom convolutional neural network specifically tailored for high-dimensional spectral data and the low cardinality of the available training set. Additionally, we address typical challenges encountered in hyperspectral analysis, including data imbalance through class balancing techniques, and interpretability via spectral band attribution analysis using Integrated Gradients. Our findings demonstrate improved interpretability and offer practical insights - for instance, the association of Sepia and Pencil with SWIR bands, and the sensitivity of Cotton and Hemp to distinct spectral regions - that can guide optimized imaging protocols and inform preservation strategies. This framework advances the non-invasive analysis of historical documents, supporting both accurate classification and interpretation for heritage conservation.

Keywords: Hyperspectral imaging, historical document analysis, spectral band attribution, cultural heritage.

1. INTRODUCTION

Historical handwritten documents are a core element of cultural heritage, as they provide precious insights into the historical, artistic, technological, and socio-economic contexts of past societies. Their preservation is of paramount importance to scholars, archivists, and conservators, yet it poses significant practical challenges, primarily due to the fragile and heterogeneous nature of their constituent materials [1]. Accurate identification and characterization of inks and substrates used in these historical manuscripts are fundamental for determining suitable preservation techniques, authenticating documents, and understanding historical practices [2].

In recent years, hyperspectral imaging (HSI) has emerged as a powerful non-destructive analytical tool in cultural heritage applications [3]. By capturing detailed spatio-spectral information across a wide wavelength range, HSI enables comprehensive characterization of document materials without physical sampling [4]. However, analyzing hyperspectral data to reliably classify inks and substrates remains challenging. This complexity arises from the high dimensionality of spectral data, the inherent variability in historical materials and their stage of conservation through the aging process, combined with frequent issues such as sensor-related noise [5] and signal degradation at the spectral extremes [6], as well as environmental and instrument-induced variability encountered in real-world acquisition scenarios [7].

Machine learning, particularly deep learning methodologies, have demonstrated substantial potential to address these challenges by learning complex patterns directly from data [8]. Despite achieving promising levels of accuracy, these models often lack transparency, offering limited insight into why specific spectral bands lead to certain classification outcomes. This lack of interpretability limits the full practical adoption of machine learning in the field of cultural heritage, where clear justifications for analytical decisions may be required. To address this, we propose leveraging attribution methods, specifically Integrated Gradients [9], to enhance the interpretability of deep learning models in hyperspectral classification tasks. Spectral attribution provides a detailed understanding of how individual bands and their nonlinear interactions contribute to ink and substrate identification. By investigating these dependencies, we aim to derive practical insights into historical material properties and inform optimized hyperspectral acquisition protocols.

In this paper, we present a detailed investigation of handwritten ink and substrate classification using hyperspectral imaging combined with a custom convolutional neural network. We implement a necessary phase of data balancing to face a common problem in cultural heritage [10], and we demonstrate the utility of spectral band attribution analysis to provide interpretability in hyperspectral classification.

2. DATASET

This experimentation involves a set of multispectral image manuscripts, acquired with Resonon PikaL and PikaIR+ line scan spectral cameras (Resonon Inc, Montana, USA [11]). The data is a subset of the HYPERDOC dataset [12] also used in derived form in [13]. More specifically, we have 66 documents, from four different sources:

- **Mock-up (25 documents)**: This set comprises synthetic samples introduced in [14], featuring historically-inspired inks (iron-gall, sepia, carbon-based inks, and their mixtures) prepared using traditional 13th- to 17th-century recipes with Arabic gum as a binder. Each ink was applied to fill a 1×1 cm square, accompanied by handwritten text. Two substrates were used: parchment and hand-crafted cotton-linen paper.

- **Provincial Historical Archive manuscripts (7 documents)**: This group includes notarial manuscripts dated between 1488 and 1494, along with one religious text of uncertain age. Ink identification, based on optical microscopy, Scanning Electron Microscopy (SEM), and Fourier-Transform Infra-Red (FTIR) spectroscopy, confirmed the presence of both pure carbon-based and iron gall inks on linen paper.

- **Genealogies book (10 documents)**: This subset features a family tree manuscript from the 16th and 17th centuries. It contains both handwritten and stamped samples using two ink types on cotton-linen paper: a mixture of iron-gall and sepia, and a pure carbon-based ink. Ink characterization was performed via SEM.

- **Royal Chancellery Archive manuscripts (24 documents)**: The final group encompasses a collection of noble lawsuit records (1459-1608) written with iron-gall ink on parchment support, along with additional samples from 18th-century maps containing iron-gall and carbon-based inks on linen and cotton paper. Ink and substrate identification relied on X-Ray Fluorescence (XRF) and optical microscopy, respectively.

Each document has been captured using the two line-scan cameras, and interpolated with a sampling interval of 5 nm:

- 121 spectral bands in the Visible and Near-Infrared (VNIR) range from 400 to 1000 nm (Pika L)
- 161 bands in the Short-Wave Infrared (SWIR) range, from 900 to 1700 nm (Pika IR+)

The VNIR hypercube has been spatially registered to align with the lower-resolution SWIR version. In this way, every document is described on a larger spectral range, from 400nm to 1700nm, with partial overlap between 900nm and 1000nm. From these 66 documents, 337 regions of interest (fragments) have been extracted and labeled in terms of the containing ink and substrate.

- Possible inks include: Iron-gall (I), Carbon (C), Sepia (S), Pencil (P). We use the term "ink" to also indicate pencil marks.
- Possible substrates include: Parchment (Pc), Cotton (Co), Linen (Li), Paper (Pp), Hemp (He).

Note that each fragment might be labelled with multiple inks and multiple substrates. For example, a mixture of Iron-gall and Carbon (I-C) is relatively common, and composite Cotton-Linen (Co-Li) paper is frequent as well. The resulting dataset is highly imbalanced in terms of class representation, an issue which will be addressed in Section 3.2.

The fragments have varying spatial resolution, ranging from the smallest side of about 60 pixels, to the largest side of about 500 pixels. Some samples are shown in Figure 1, with false colors extracted from the VNIR range.

Figure 1. Sample fragments used within this study, visualized in false RGB colors from the VNIR range (630, 535 and 430 nm bands).

For our experiments, the documents were partitioned into two sets: 49 documents for train (245 fragments), and 17 documents for test (92 fragments), thus resulting in a 73% - 27% split over the fragments. Partitioning at document level ensures that no two fragments from the same document are split between training and test sets. Such precaution prevents the introduction of bias in the training-test split, which would create an unrealistic evaluation scenario, where test performance does not reflect real-world performance.

3. METHODOLOGY

By the very nature of the involved data, we are dealing with two multilabel classification tasks: one related to ink classification, and one related to substrate classification. An alternative approach to the problem might have co-occurrences of different inks and different substrates gathered into joint classes (e.g. creating distinct classes for "Iron-gall", "Carbon", and "Iron-gall/Carbon mixture"). We reserve for future work this line of research. In the following, we describe the neural architecture that was designed to address the two multilabel classification tasks, as well as our procedure for class balancing.

3.1 Neural Network

We design and present a custom neural network to address the problem of ink and substrate multilabel spectral classification, as displayed in Table 1.

Table 1. Architecture of the custom neural network designed for ink and substrate multispectral classification. Indicated settings include output size (o), kernel size (k), stride (s), dropout probability (p).

Layer	Settings	Input shape	Output shape	# Params
Input		(128, 128, 282)	(128, 128, 282)	0
Conv	o=128 k=3×3 s=2	(128, 128, 282)	(64, 64, 128)	324,864
BatchNorm		(64, 64, 128)	(64, 64, 128)	256
ReLU		(64, 64, 128)	(64, 64, 128)	0
MaxPool	k=3×3, s=3	(64, 64, 128)	(22, 22, 128)	0
Dropout	p=0.3	(22, 22, 128)	(22, 22, 128)	0
Conv	o=64, k=3×3, s=2	(22, 22, 128)	(11, 11, 64)	73,728
BatchNorm		(11, 11, 64)	(11, 11, 64)	128
ReLU		(11, 11, 64)	(11, 11, 64)	0
MaxPool	k=3×3, s=3	(11, 11, 64)	(4, 4, 64)	0
Dropout	p=0.3	(4, 4, 64)	(4, 4, 64)	0
Conv	o=32, k=3×3, s=2	(4, 4, 64)	(2, 2, 32)	18,432
BatchNorm		(2, 2, 32)	(2, 2, 32)	64
ReLU		(2, 2, 32)	(2, 2, 32)	0
GlobalMaxP.	o=1×1	(2, 2, 32)	(1, 1, 32)	0
Flatten		(1, 1, 32)	(32)	0
Linear	o=16	(32)	(16)	528
Dropout	p=0.5	(16)	(16)	0
Linear (Ink)	o=4	(16)	(4)	68
Linear (Substrate)	o=5	(16)	(5)	85

The model takes as input spectral matrices shaped (128×128×282): the 128×128 spatial resolution is obtained via random cropping, padding the fragment whenever needed (i.e. if the initial resolution is smaller than 128 pixels for any side). The 282 channels are obtained as a concatenation of VNIR (121) and SWIR (161) information. Our architecture is structured with three consecutive convolutional blocks, including batch normalization, rectifying linear unit (ReLU) activation, local max pooling, and dropout. These are followed by a global max pooling and a linear layer mapping to 16 abstract features. From this stage, two parallel linear branches are built respectively for ink classification (with 4 possible

outputs) and substrate classification (with 5 possible outputs). Consequently, two loss functions based on binary cross-entropy with logits are used to compare the outputs with the available multilabel ground truth, and guide gradient backpropagation for training. At inference time, the loss functions are replaced with sigmoid functions to constrain the output into independent probability-like values in the [0-1] range: for any given input, different ink classes and different substrate classes are associated with output probabilities, which are then thresholded to 0.5 (50%) for hard decisioning.

We train our model for 250 epochs, which are found to be sufficient to reach model convergence. We adapt model weights through the Adam optimizer, with learning rate empirically set to 0.05. We select the highest possible batch size to 16 due to hardware constraints. Details on the hardware and software configuration are provided in Section 4.1.

Note that, through preliminary experiments, we also evaluated the effectiveness of existing model families for image classification, such as ResNet (versions 18, 34, 50), EfficientNet (versions b0, b1, b2, b3), and MobileNet (versions v2, v3small, v3large). These solutions were pretrained on the IMAGENET classification task, which focuses on color images depicting animals and objects. The radically different nature of our task (high-dimensional spectral images of historical documents), together with the low cardinality of the available training set, and the relatively high complexity of said architectures, resulted in extremely poor performance, which led us to eventually focus on our custom solution.

3.2 Class Balancing

The training set cardinality of the dataset used is reported on the left side of Table 2. Relative distribution is reported in parentheses, to highlight the situation of class unbalancing, as also visualized in the left side of Figure 2 (a) and (b).

We opt for super-sampling as a technique to address the class unbalancing, whose effect on cardinality is shown on the center columns of Table 2 and on the right side of Figure 2 (a) and (b). Note that, given the multilabel nature of our classification tasks, replicating one sample will, in general, impact multiple classes across both class sets. For this reason, it is not possible to produce a perfectly uniform distribution of classes. We will quantify the quality of our handcrafted super-sampling scheme in Section 4.2.

Table 2. Classes cardinality for the training and test sets of Inks and Substrates class sets. For the training set, we report cardinalities before and after balancing. Relative distribution reported in parentheses.

Class set	Classes	Cardinality (train)		Balanced cardinality (train)		Cardinality (test)	
Inks	Iron-gall	136	(27,9%)	200	(20,9%)	52	(33.3%)
	Carbon	100	(20,5%)	217	(22,6%)	26	(16.7%)
	Sepia	67	(13,7%)	213	(22,2%)	22	(14.1%)
	Pencil	185	(37,9%)	329	(34,3%)	56	(35.9%)
Substrates	Parchment	55	(16,0%)	74	(13,2%)	9	(8.2%)
	Cotton	76	(22,1%)	113	(20,1%)	18	(16.4%)
	Linen	151	(43,9%)	207	(36,9%)	28	(25.5%)
	Paper	22	(6,4%)	113	(20,1%)	18	(16.4%)
	Hemp	40	(11,6%)	54	(9,6%)	37	(33.6%)

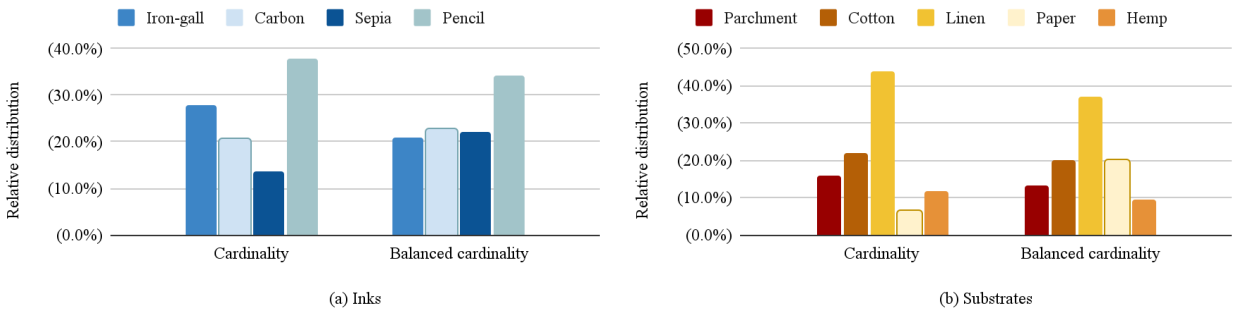

Figure 2. Relative classes cardinality for the training set of Inks and Substrates class sets, before and after balancing.

3.3 Spectral Band Attribution

We set out to study the impact of each input spectral band on the model's prediction, starting from two considerations:

1. Some bands may be useful for material classification only in conjunction with others (e.g., ratios or differences between specific wavelengths may be as important as their individual values).

2. Nonlinear interactions in the neural model could amplify or suppress the importance of certain bands based on their combined effect.

Consequently, we adopt Integrated Gradients [9]: a method for input-output attribution that captures nonlinear interactions and that is not limited to addressing the model's behavior at a specific point in the input space. Let x' be a baseline input: an all-zeros reference spectral image that represents the absence of features. Let x be an actual input, whose attribution over output T of model F we want to describe. Let α be a scaling factor that interpolates between x' and x.

$$I(x) = (x - x') \int_{\alpha=0}^{1} \frac{\partial F_T(x' + \alpha(x - x'))}{\partial x} d\alpha, \quad (1)$$

i.e. the attribution I(x) is computed integrating the gradients of the model's prediction output class T with respect to the input along a straight line path from the baseline to the actual input. The result is a matrix with the same spatial and spectral dimensions as the input x. To obtain the attribution per band, we then sum over the spatial dimensions, and obtain a vector with our number of spectral bands. We run this for all images, analyzing attribution information with respect to the predicted classes T. Observations are presented in Section 4.4.

4. EXPERIMENTAL RESULTS

4.1 Hardware and Software Configuration

Our solution is implemented in Python 3.10.12 using the PyTorch 2.3.0 framework. We run our experiments on a machine with the following configuration:

- Processor (CPU): Intel(R) Core(TM) i7-7700 CPU @ 3.60GHz (8 CPUs)
- Memory (RAM): 32 GB
- Graphics Card (GPU): NVIDIA Titan X, 12 GB
- Storage: 3 TB ext4 SSD
- Operating System: Ubuntu 22.04.3 LTS, 64-bit.

4.2 Balancing Evaluation

In this section we evaluate the effect of class balancing using metrics for the assessment of distributions. The impact on classification itself is evaluated in Section 4.3. We select the following metrics:

Entropy: $H = -\sum_{i=1}^{N} p_i \, log\,(p_i)$ (2)

Normalized entropy: $H' = \frac{H}{log(N)}$ (3)

Gini coefficient: $G = 1 - \sum_{i=1}^{N} p_i^2$ (4)

Class imbalance ratio: $CIR = \frac{min(n_i)}{max(n_i)}$ (5)

where n_i is the count of instances in class i, N is the total number of classes, and $p_i = \frac{n_i}{\sum_{j=1}^{N} n_j}$.

The results, reported in Table 3, clearly show an improvement for all metrics as given by our super-sampling scheme for class balancing.

Table 3. Class balancing metrics as evaluated on the training set in its unbalanced and balanced version. For all metrics, the higher the better. Best results per class set in boldface.

Class set	Balancing	H ↑	H' ↑	G ↑	CIR ↑
Inks	Unbalanced	1.3212	0.9531	0.7178	0.3622
	Balanced	**1.3644**	**0.9842**	**0.7383**	**0.6079**
Substrates	Unbalanced	1.4142	0.8787	0.7153	0.1457
	Balanced	**1.5059**	**0.9357**	**0.7560**	**0.2609**

4.3 Classification Results

We here report aggregate classification metrics on the test set. The results, shown in Table 4, are expressed in terms of precision, recall, and F1 score. The underlying model is trained either on class-unbalanced or class-balanced data, while the test set is kept in its original unbalanced form. Consequently, we report macro statistics computed as the average of class-specific metrics, in addition to global micro statistics. The aggregate results show a consistent positive impact of class balancing on classification performance, demonstrated by improved statistics across almost all metrics.

Table 4. Aggregate classification results on the test set, with models trained respectively on class-unbalanced and class-balanced data. For all metrics, the higher the better. Best results per class set in boldface.

Class set	Balancing	Precision ↑		Recall ↑		F1 score ↑	
		Micro	Macro	Micro	Macro	Micro	Macro
Inks	Unbalanced	**75%**	60%	71%	61%	**73%**	59%
	Balanced	72%	**72%**	**74%**	**74%**	**73%**	**72%**
Substrates	Unbalanced	82%	54%	**65%**	53%	72%	57%
	Balanced	**89%**	**89%**	64%	**61%**	**74%**	**67%**

Detailed per-class classification results are consequently reported in Table 5, related to the class-balanced experiment. Note that the multilabel nature of the task does not allow the construction of a confusion matrix. Although class balancing did improve the overall model behavior, it can still be observed how the classes with lowest F1 score are also the classes with originally lowest cardinality, i.e.: Sepia for the Inks class set, and Paper for the Substrates class set. This highlights the importance of collecting new real data as a direction for further improvement.

Table 5. Detailed classification results on the test set, with models trained on class-balanced data. For all metrics, the higher the better.

Class set	Class	Precision ↑	Recall ↑	F1 score ↑
Inks	Iron-gall	65%	69%	67%
	Carbon	96%	96%	96%
	Sepia	34%	55%	42%
	Pencil	93%	77%	84%
Substrates	Parchment	78%	78%	78%
	Cotton	88%	38%	54%
	Linen	85%	100%	92%
	Paper	100%	22%	36%
	Hemp	96%	65%	77%

4.4 Spectral Band Attribution

In Figure 3 and Figure 4 we present the results of spectral band attribution based on the Integrated Gradients method, respectively on ink classes and substrate classes.

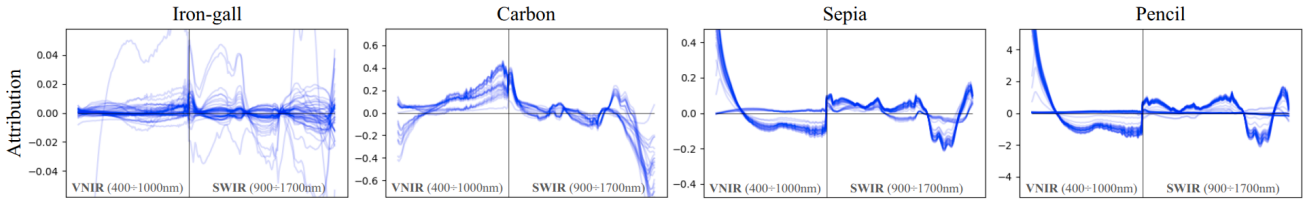

Figure 3. Integrated Gradients attribution for VNIR and SWIR bands on ink classes. The positive, or negative, attribution is correlated directly, or inversely, with assigning the given class.

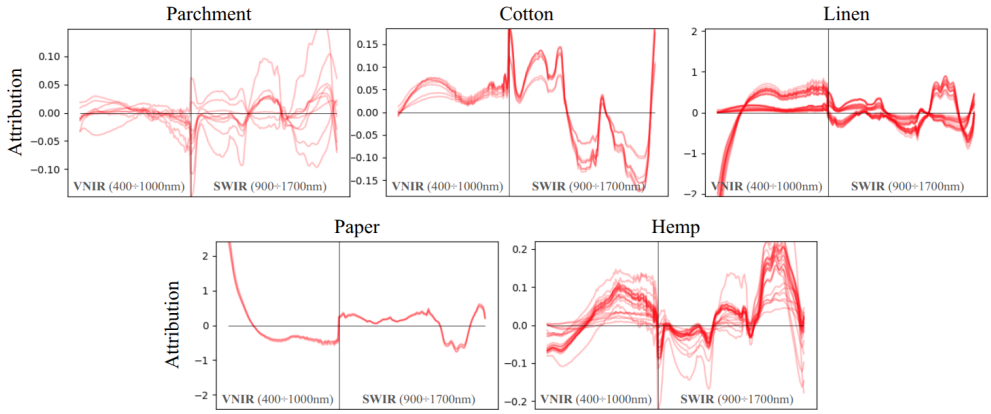

Figure 4. Integrated Gradients attribution for VNIR and SWIR bands on substrate classes. The positive, or negative, attribution is correlated directly, or inversely, with assigning the given class.

Every individual plot represents one class, with every line representing one image that was predicted as belonging to that class. The horizontal axis covers the 121 VNIR bands, followed by the 161 SWIR bands. The vertical axis indicates attribution: a positive attribution value indicates that the given spectral band contributes to identifying the class, whereas a negative attribution indicates that the band contributes to reducing the confidence on the class. A null attribution indicates that the band does not influence that specific class.

Concerning inks, Iron-gall (I) is given relatively equal contribution from all spectral bands. Carbon (C) displays negative contribution from the spectral extremes, where data is known to be noisy [6], yet the model is not learning to ignore it; a higher contribution from VNIR bands compared to SWIR bands is also observed. Sepia (S) and Pencil (P) exhibit a reciprocally-similar behavior: signal presence in the SWIR bands functions as a strong marker for the class, whereas signal in the VNIR bands has the opposite effect; spectral extrema appear to have a strong influence in class prediction, which suggests possible overfitting to the training data, given the known noisy nature of the involved spectral ranges. Concerning substrates, Parchment (Pa) tends to focus more on the SWIR range. Cotton (Co) exhibits a negative correlation with respect to higher wavelengths (second half of the SWIR bands). Linen (Li) clearly shows two different behaviors, which are to be attributed to the two types of images effectively contained within this class: pure linen, and cotton-linen. Finally, Hemp (He) presents a high correlation with high VNIR wavelengths and high SWIR wavelengths.

In general terms, a procedure of spectral band attribution is particularly useful if paired with domain knowledge, in order to gain insights and draw conclusions concerning specific behaviors of the model in relation to the data. A significant example of this is the mentioned role of noisy spectral extrema in the model's prediction.

5. CONCLUSIONS

In this study, we demonstrated the effectiveness of hyperspectral imaging combined with deep learning for accurately classifying inks and substrates in historical handwritten documents. Our proposed custom convolutional neural network

successfully handled the inherent complexity of high-dimensional spectral data and low training cardinality. Addressing data imbalance through class balancing, then, further improved classification performance.

Thanks to spectral band attribution using Integrated Gradients, we were able to inspect the underlying connections between specific bands and output classes within our model. Sepia and Pencil inks, for example, showed strong positive attribution in the SWIR range, with negative contributions from VNIR bands, suggesting that their detection is driven primarily by longer wavelengths. Conversely, Carbon ink presented complex, mixed interactions across the VNIR range, indicating a more diffuse spectral signature. This behavior aligns with the known characteristic of carbon-based inks exhibiting consistently flat, low reflectance across both VNIR and SWIR wavelengths: to reliably detect this flat spectral profile, the model must analyze both spectral ranges. Among substrates, Cotton showed negative correlation with higher SWIR wavelengths, while Hemp was strongly associated with both upper VNIR and SWIR bands. These patterns are in part confirmed by known spectral behaviors, but they also highlight model-specific behaviors that could form the basis for further insights into spectral-based analysis of historical documents.

Future work may extend spectral attribution to broader datasets and explore its combination with domain-specific knowledge to assist preservation experts in targeted material analysis and imaging protocol design [15].

REFERENCES

[1] S. Zervos and I. Alexopoulou, "Paper conservation methods: a literature review," Cellulose 22, 2859–2897, (2015).

[2] K. Nesměrák and I. Němcová, "Dating of historical manuscripts using spectrometric methods: a mini-review," Analytical Letters 45, 330–344, (2012).

[3] C. Balas, V. Papadakis, N. Papadakis, A. Papadakis, E. Vazgiouraki and G. Themelis, "A novel hyper-spectral imaging apparatus for the non-destructive analysis of objects of artistic and historic value," Journal of Cultural Heritage 4, 330–337, (2003).

[4] R. Padoan, T.A. Steemers, M. Klein, B. Aalderink and G. De Bruin, "Quantitative hyperspectral imaging of historical documents: technique and applications," Art Proceedings, 25–30, (2008).

[5] B. Rasti, P. Scheunders, P. Ghamisi, G. Licciardi and J. Chanussot, "Noise reduction in hyperspectral imagery: Overview and application," Remote Sens. 10, 482, (2018).

[6] L. Zhuang, L. Gao, B. Zhang, X. Fu and J.M. Bioucas-Dias, "Hyperspectral image denoising and anomaly detection based on low-rank and sparse representations," IEEE Trans. Geosci. Remote Sens. 60, 1–17, (2020).

[7] R.L. Wiggins, L.E. Comstock and J.J. Santman, "Real-world noise in hyperspectral imaging systems," in Airborne Intelligence, Surveillance, Reconnaissance (ISR) Systems and Applications VIII, vol. 8020, 31–37, SPIE, (2011).

[8] S. Li, W. Song, L. Fang, Y. Chen, P. Ghamisi and J.A. Benediktsson, "Deep learning for hyperspectral image classification: An overview," IEEE Trans. Geosci. Remote Sens. 57, 6690–6709, (2019).

[9] M. Sundararajan, A. Taly and Q. Yan, "Axiomatic attribution for deep networks," in Int. Conf. Mach. Learn., 3319–3328, PMLR, (2017).

[10] Rei, D. Mladenic, M. Dorozynski, F. Rottensteiner, T. Schleider, R. Troncy, J.S. Lozano and M. Gaitán Salvatella, "Multimodal metadata assignment for cultural heritage artifacts," Multimedia Syst. 29, 847–869, (2023).

[11] Resonon Inc., "Resonon Hyperspectral Imaging Cameras," Resonon website, accessed July 2025, available: https://resonon.com/hyperspectral-cameras

[12] A.B. López-Baldomero, J.L. Nieves, F. Moronta-Montero, M.Á. Martínez-Domingo, R. Fernández-Gualda, J. Hernández-Andrés, A.S. Reichert et al., "Hyperspectral dataset of historical documents and mock-ups from 400 to 1700 nm (HYPERDOC)," Scientific Data 12, 1248, (2025).

[13] A.B. López-Baldomero, M. Buzzelli, F. Moronta-Montero, M.Á. Martínez-Domingo and E.M. Valero, "Ink classification in historical documents using hyperspectral imaging and machine learning methods," Spectrochim. Acta A Mol. Biomol. Spectrosc. 335, 125916, (2025).

[14] A.B. López-Baldomero, E.M. Valero, A.S. Reichert, F. Moronta-Montero, M.Á. Martínez-Domingo and A. López-Montes, "Hyperspectral database of synthetic historical inks," in Archiving Conf., 11–16, (2024).

[15] I.N. Tzortzis, I. Rallis, K. Makantasis, A. Doulamis, N. Doulamis and A. Voulodimos, "Automatic inspection of cultural monuments using deep and tensor-based learning on hyperspectral imagery," in ICIP, 3136–3140, (2022).