Singh *et al. Discover Computing* (2025) 28:333
https://doi.org/10.1007/s10791-025-09837-2

Discover Computing

RESEARCH Open Access

Unsupervised deep autoencoder-based reconstruction for ink mismatch detection in hyperspectral document images

Pangambam Sendash Singh[1], Subbiah Karthikeyan[2], Govind Murari Upadhayay[3], Sounak Sadhukhan[4], Pramod Kumar Soni[3*] and Timothy Malche[3]

*Correspondence:
Pramod Kumar Soni
pramod.soni@jaipur.manipal.edu
[1]Department of Computer Science and Applications, Doctor Harisingh Gour Vishwavidyalaya, Sagar, Madhya Pradesh 470003, India
[2]Department of Computer Science, Banaras Hindu University, Varanasi, Uttar Pradesh 221005, India
[3]Department of Computer Applications, Manipal University Jaipur, Jaipur, Rajasthan 303007, India
[4]School of Computer Science Engineering and Technology, Bennett University, Greater Noida, Uttar Pradesh 201310, India

Abstract
Ink analysis is crucial for finding ink mismatches in handwritten documents to determine their authenticity and identify potential forgery. Traditional chemical-based methods like thin layer chromatography are effective for this purpose, but they are destructive, irreversible, time-consuming, and sensitive to environmental factors. Hyperspectral document images (HSDIs) which capture the spectral information in multiple bands can reveal the material composition, thereby allowing the identification of different inks by their distinct spectral characteristics, even when the inks visually appear to be the same colour. Hyperspectral imaging thus shows great promise for forensic document analysis and authentication. Despite its potential, research in this area is still emerging. Existing HSDI-based methods show promise in detecting ink mismatches but these methods often require prior spectral information and ground truth data for the ink pixels, limiting their practical applicability. Unsupervised methods present a solution by removing the need for prior information. This work proposes a novel unsupervised ink mismatch detection method in HSDIs using deep learning-based hyperspectral analysis. The proposed method formulates ink mismatch detection as a decision problem that determines whether a deep autoencoder trained on the spectral features of specific ink pixels would be able to reconstruct the spectral features of unseen complementary ink pixels. The proposed framework works in a fully unsupervised manner, learning the spectral representations directly from data without any labeled samples or prior spectral information of the inks present. Unlike most existing methods which are supervised, the proposed unsupervised method is more practically applicable for real world forensic document analysis. Experimental results on a publicly available HSDI dataset demonstrate the superiority of the proposed method over existing methods in identifying ink mismatches in potentially fraudulent document images.

Keywords Hyperspectral document analysis, Deep learning, Autoencoder, Forgery detection, Ink mismatch detection, Industry innovation

1 Introduction

Forensic document analysis aims to address the question of authenticity of documents suspected of being fraudulent through scientific examination of any unscrupulous manipulation made to them. Human eyes can perceive a wide range of colours in the visible spectrum; nevertheless, they face difficulties in distinguishing between two apparently similar colours. Fraudulent documents often involve modifications made with ink colours that closely resemble to those used elsewhere in the document, making it challenging for visual differentiation with the human eye. For this reason, ink analysis for detection of ink mismatch becomes crucial in determining document forgery. Hyperspectral images [1] that use both spectroscopy and imaging technology concepts to capture objects in multiple densely sampled wavelengths have better discrimination and characterization capabilities than regular visual images [2, 3]. This technology, although originally developed for remote sensing based applications, has recently gained attention as a promising tool for analyzing and authenticating forensic traces [4, 5].

Spectral based methods have been utilized in various forensic applications, including counterfeit detection [6, 7], blood-contaminated fingermark detection [8, 9], and blood stain age estimation [10], etc. Early chemical analysis based ink mismatch detection methods, such as thin layer chromatography [11], which involved separating ink into its colour components using chemical analysis, are effective; however these chemical approaches are destructive, irreversible, time-consuming, and sensitive to environmental factors, and hence they require a lot of precautions. On the other hand, nondestructive spectral imaging techniques can detect even subtle differences in ink colours through hyperspectral analysis of ink pixels, thereby offering a nondestructive alternative for ink mismatch detection in forensic documents.

Hyperspectral imaging based conventional ink mismatch detection methods often start with manual identification of suspected areas within a document. Such manual approaches are both time-consuming and error-prone, thereby highlighting the need for automated approaches. While the research in this field is still relatively new, existing methods have demonstrated promising results for detecting ink mismatches in hyperspectral document images (HSDIs) [12, 13]. Nevertheless, most of these existing methods are often limited by the requirement for prior spectral knowledge and ground truth data for the ink pixels thereby restricting their practical applicability in real-world scenarios. Unsupervised methods provide an alternate solution to address this challenge by eliminating the need for prior information, thereby offering a more practical method for detecting ink mismatches in forensic document analysis.

In this work, an unsupervised deep spectral reconstruction based method is proposed for automated ink mismatch detection in HSDIs. The proposed method formulates the task of ink mismatch detection as a decision problem that determines whether a locally trained autoencoder would effectively be able to reconstruct spectral features for an unseen complementary set of ink pixels. Despite the growing interest in spectral analysis for forensic applications, only a few studies have explored unsupervised approaches in this domain. We summarize this work's major contributions as follows:

1. Introduction of an efficient unsupervised ink mismatch detection method for HSDIs through autoencoder based spectral reconstruction.

2. Consideration of both local and global relationships among ink pixels, offering a comprehensive solution for detecting subtle ink mismatches in forensic document.
3. Demonstrating the effectiveness and practical applicability of the proposed method through experimentation and validation with real-world data.
4. Opening up possibilities for further research in unsupervised ink mismatch detection methods to achieve more robust and accurate forensic document analysis.

The remainder of the paper is organized as follows: Sect. 2 discusses prior research related to ink mismatch detection. Section 3 outlines the proposed method along with a brief introduction to autoencoders. The experimental setup and result analysis are presented in Sect. 4, followed by an ablation study in Sect. 5. Finally, Sect. 6 presents the conclusions of the study along with future research directions.

2 Previous studies

Despite the significant potential of hyperspectral imaging (HSI) in addressing complex challenges in document image analysis, hyperspectral document image (HSDI) research remains relatively limited. Several works have discussed its applications, challenges, and future prospects in document forensics [14–16]. The research on ink mismatch detection has evolved from traditional spectral analysis techniques to advanced machine learning and deep learning-based reconstruction frameworks.

2.1 Traditional and statistical approaches

Early studies focused on physics-based and statistical methods for ink differentiation. In [17], a Fourier transform spectroscopy system combined with Fuzzy C-means clustering was used for differentiating inks, though it was limited by slow image acquisition and qualitative analysis. The work in [18] used near-infrared (NIR) imaging with chemometric techniques for identifying document forgeries, while [19] applied Least Squares Support Vector Machines (LS-SVM) to bank cheque forgery detection using hyperspectral data.

2.1.1 Clustering and feature-based methods

Several clustering approaches were developed for ink discrimination in hyperspectral documents. The works in [20, 21] utilized spectral partitioning of ink pixels, whereas [22] proposed clustering and outlier detection to estimate the number of unique inks in localized document regions. In [5], a sequential forward band selection method that aims to improve detection accuracy is proposed. The limitation for the proposed method is its requirement of prior information about ink proportions. Hyperspectral unmixing techniques [23] demonstrated improved performance and robustness against unbalanced ink clusters in estimating spectral signatures and ink proportions.

2.1.2 Machine learning and deep learning based approaches

With the increasing availability of hyperspectral data, machine learning models became prominent for ink analysis. In [24], forgery detection in legal documents was formulated as a binary classification problem using statistical features and classifiers such as Multilayer Perceptron, RBF-SVM, Random Forest, and CNNs. The study by [25] applied CNNs for writer identification using spectral features extracted from hyperspectral text regions. Recent advancements have explored deep learning frameworks for ink mismatch detection. Convolutional Neural Networks (CNNs) using spectral [13] and combined spectral-spatial information [26] achieved improved accuracy but required prior spectral knowledge, limiting their real-world applicability. Autoencoder-based approaches have shown promise for unsupervised feature learning [27] employed Convolutional Autoencoder (CAE)-based features with Logistic Regression for ink discrimination, while [28] used CAE-derived features with SVM for writer identification. Similarly, [29] introduced a graph-based orthogonal Nonnegative Matrix Factorization (NMF) method for unsupervised ink forgery detection. Furthermore, recent studies [30, 31] utilized Convolutional Autoencoder features combined with Logistic Regression for ink mismatch detection. These approaches achieved high accuracies, however they rely on labeled data and prior knowledge of ink types, which limits their applicability in real-world forensic scenarios where the inks used may be unknown.

2.1.3 Recent developments and hybrid models

More recent efforts have combined hyperspectral unmixing and clustering for improved ink discrimination. The studies presented in [32, 33] integrated k-means and Gaussian Mixture Models to handle unbalanced ink clusters and enhance entropy-based ink unmixing. In another direction, [34] proposed combining ion mobility spectrometry (IMS) with machine learning classifiers to non-destructively classify inks from various manufacturers with high accuracy. Although substantial progress has been made, the approaches available in the literature often rely on prior spectral knowledge of the inks, which limits their practical applicability. Additionally, deep learning based methods employing CNNs for hyperspectral document analysis suffer from slow processing speeds, thereby restricting such methods from real-time or large-scale applications. Moreover, the generalization ability and robustness of such methods in different document types, ink compositions, and environmental conditions still remain uncertain.

Forgery detection has also been studied in natural and digital document images using RGB-based features. In [35] a differential abnormality-based tampering detection method is proposed. It explores the local variances in RGB intensity and structure to identify forged regions in a document. In [36], the authors improve localization accuracy in complex manipulated images by using semantic and chromatic cues information. In [37], a robust multi-view contrastive learning approach is used to distinguish authentic and tampered regions in a document image. The authors in [38] utilize transformer based self-attention mechanisms to model long-range dependencies to achieve fine-grained localization of tampered areas in a manipulated document. These RGB-based methods explore visual and semantic features to identify document forgeries. The proposed method in this study, on the other hand, diverges by exploring hyperspectral information to detect ink-level discrepancies, where discrimination arises from spectral rather than visual context.

In summary, document forgery detection has evolved in four main ways: (i) traditional and statistical methods that use physics-based analysis, (ii) clustering and feature-based methods that focus on spectral partitioning and unmixing, (iii) machine learning and deep learning models that automate feature extraction and classification, and (iv) recent hybrid and RGB-based forgery detection strategies that focus on spatial and semantic localization. RGB-based methods work well for semantic forgery localization, but hyperspectral imaging is the only method that can tell the difference between ink compositions, which is crucial for forensic document examination. The proposed deep autoencoder-based framework in this work is built on this foundation by using unsupervised spectral feature learning to improve the detection of ink-level mismatches in HSDIs.

3 Materials and method

This section outlines the methodology employed in this study. It begins with a description of the autoencoder architecture and its specific role in the proposed approach. Subsequently, the details of the proposed method are presented.

3.1 Autoencoder

An autoencoder [39, 40] is a special kind of feed-forward neural network that contains an input layer, a hidden/code layer, and an output/reconstruction layer [41]. The basic architecture of an autoencoder is shown in Fig. 1 (left). In an autoencoder, the number of nodes in the output layer is same as the number of nodes in the input layer. During its training, the autoencoder maps the input $x \in \mathbb{R}^d$ to the output layer, reconstructing x through a lower-dimensional latent space by using activation functions between layers. The portion of the network that maps (encodes) the higher dimensional input to a lower dimension feature vector (y) is known as *encoder*. Further, the remaining portion of the network called the *decoder* takes this lower dimensional feature vector (y) as an input and maps it to the reconstruction layer to recreate $\hat{x} \in \mathbb{R}^d$, of the input. These mapping functions can be expressed mathematically as follows:

$$y = f\left(W_x x + \beta\right), \quad \hat{x} = f\left(W_y y + \gamma\right) \tag{1}$$

where, W_x and W_y are the input to hidden layer and from hidden to reconstruction layers' weights respectively, β and γ represent the bias of the hidden and reconstruction layers, $f(.)$ is the non-linear activation function that is used in between the layers. For an input training dataset $X = \{x_1, \ldots, x_N\}, \; x_i \in \mathbb{R}^d$, the objective of autoencoder training using X is to minimize the reconstruction error,

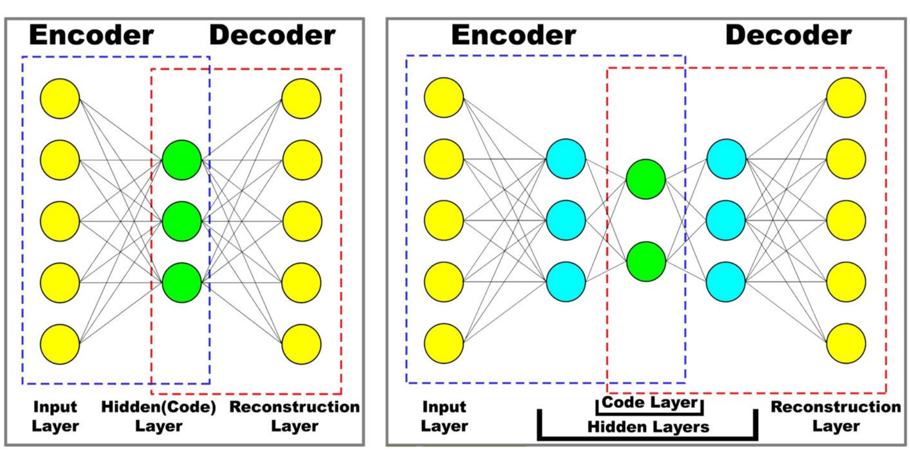

Fig. 1 Network architecture of a simple and a stacked autoencoder

$$Minimize : \sum_{i=1}^{N} \|x_i - \hat{x}_i\|^2 \tag{2}$$

The primary objective of an autoencoder is to learn a compact feature representation (y) of the input vectors (x) satisfying Eqs. 1 and 2. This process involves iterative updations of W_x and W_y is such a way that the reconstruction error decreases gradually in every consecutive iteration until it is below some threshold. In order to reduce reconstruction errors, simple autoencoders are often arranged such that the hidden layer of one serves as the input for the next, creating a "Stacked Autoencoder" (SAE) [40] (Fig. 1, right).

3.2 Proposed method

The proposed method for detecting ink mismatches relies on spectral reconstruction [42, 43] achieved through deep autoencoders. Autoencoders, which are trained through data-driven processes, aim to learn approximated identity mappings, which are defined in terms of the interconnection weights between layers. Given the nature of its training, a well-trained autoencoder is anticipated to exhibit higher reconstruction errors when presented with data points whose distributions differ from those in the original training dataset. The proposed unsupervised ink mismatch detection approach for HsDIs is devised by leveraging this principle.

To provide a clearer overview of the workflow, this section elaborates on each stage of the proposed unsupervised ink mismatch detection pipeline from preprocessing to clustering. Let $X \in \mathbb{R}^{m \times n \times d}$ be the given HSDI, where (m, n) corresponds to the spatial size of X and d, the number of spectral bands in X. The proposed method proceeds through the following steps:

1. **Preprocessing and Segmentation:**

   - Segmentation of ink pixels from the background: Sauvola's local thresholding method [44] is utilized for segmentation of ink pixels. Such local thresholding based segmentation is deemed better than global thresholding methods like Otsu's [45] due to the variable illuminations across the document.
   - Output: A binary mask $M \in \{0, 1\}^{m \times n}$ enabling the separation of ink pixels from the background [20].

2. **Spatial Partitioning:**

   - The segmented document is divided into two disjoint spatial groups, G_1 and G_2, using sliding windows [5] of a predefined size. Each group consists of spectrally similar yet spatially distinct ink pixels. The aim of this complementary grouping is to allow each autoencoder to learn the spectral features from one region and be evaluated in another for unbiased detection of spectral variations in unseen ink pixels.

   The partitioning into two disjoint spatial segments is based on real-life document manipulations that often involve different ink types across distinct continuous spatial areas.

3. **Training of Autoencoders:** Next, the proposed framework proceeds with an unsupervised learning approach, where autoencoders are trained on spectral features of their respective ink pixels without using any labels or prior class information. This is what makes the proposed method useful in real-world forensic applications where we seldom have ground truth data.

    - Two identical deep autoencoders, AE_1 and AE_2, are trained exclusively with spectral features of the ink pixels from G_1 and G_2, respectively.
    - The goal of each autoencoder is to learn neural network-based approximated identity mappings that are specific to the spectral characteristics of its respective group.

This unsupervised training allows the network to learn intrinsic spectral representations of inks without needing class labels. This makes the network effective at finding spectral deviations that could mean potential ink mismatches. In this study, two deep autoencoders are trained separately on spectral features from different spatial regions to more accurately capture subtle spectral variations between inks. Figure 1 illustrates the generic multilayer structure of a deep autoencoder for conceptual clarity, not the two-network configuration used in this work.

4. **Evaluation of Autoencoder Reconstruction:** Upon training, the reconstruction capabilities of AE_1 and AE_2 are assessed by testing them on the pixel points from their complementary groups as follows:

    - The reconstruction errors (ε_1 and ε_2) are calculated for each autoencoder while reconstructing pixel points belonging to their trained groups.
    - The reconstruction errors ($\hat{\varepsilon_1}$ and $\hat{\varepsilon_2}$) for each autoencoder while reconstructing pixel points from the complementary groups are also computed.
    - Next, the reconstruction probabilities of each autoencoder in the complementary groups (RP_1 and RP_2) are calculated. The reconstruction probability for each pixel is computed by transforming its reconstruction error into a probabilistic confidence score using a Gaussian kernel function, as defined in [46]. This conversion reflects the likelihood that a pixel's spectral feature conforms to the learned data distribution of the corresponding autoencoder.

5. **Ink Mismatch Detection:**

    - The reconstruction probabilities are first compared with predefined threshold estimates (θ_1 and θ_2) to obtain a baseline discrimination result. These preliminary thresholds serve as starting points for further refinement using the automated optimization procedure described in Sect. 5.2.
    - If $RP_1 \geq \theta_2$ or $RP_2 \geq \theta_1$, at least one autoencoder is capable of reconstructing the spectral features of the ink pixels belonging to both groups; concluding no ink mismatch in the questioned document.

- If $RP_1 < \theta_2$ and $RP_2 < \theta_1$, none of the autoencoders are able to reconstruct the spectral features of the inks pixels from the complementary groups, inferring the presence of ink mismatch in the document.

6. **Clustering for Categorization:** In case of detecting an ink mismatch, isolate the ink pixels into two categories using K-means clustering.

   - Let N be the total number of ink pixels in the questioned document.
   - Define $\mathcal{E}_1 \in \mathbb{R}^N$ and $\mathcal{E}_2 \in \mathbb{R}^N$ as the reconstruction errors of AE_1 and AE_2 respectively, while reconstructing the spectral features of all the ink pixels.
   - Use K-means clustering with $(K = 2)$ to group all N ink pixels into two categories that correspond to different ink types. The clustering utilizes $\mathcal{E}_1$ and $\mathcal{E}_2$ as input features, with initialization performed using the K-means++ strategy. The algorithm iterates until cluster assignments stabilize or the maximum iteration limit is reached, whichever is earlier.

The core mechanism of the proposed method lies in reconstruction capability of the autoencoders that learns the spectral distribution of ink pixels that are inside their respective training regions. By using reconstruction probabilities from two complementary autoencoders, the method distinguishes spectrally consistent (meaning same ink) and inconsistent (meaning different inks) regions. Figure 2 presents the block diagram of the proposed ink mismatch detection method for HSDIs. A corresponding pseudocode representation is also given in Algorithm 1 to enhance procedural clarity and reproducibility.

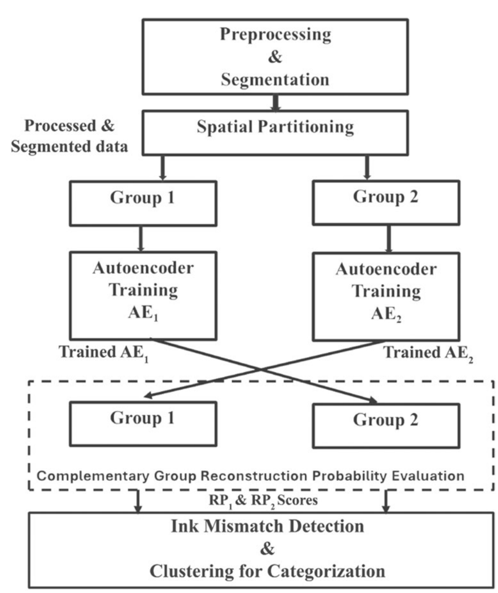

Fig. 2 Block diagram of the proposed ink mismatch detection approach

Require: HSDI $X \in \mathbb{R}^{m \times n \times d}$
Ensure: Ink mismatch detection result and categorized ink regions
Step 1: Preprocessing and Segmentation
Apply Sauvola's local thresholding on X to obtain binary mask M
Extract ink pixels: $X_{ink} = X \odot M$
Step 2: Spatial Partitioning
Partition X_{ink} into two disjoint spatial groups G_1 and G_2 using sliding window
Step 3: Training Autoencoders
Train AE_1 using spectral features from G_1
Train AE_2 using spectral features from G_2
Step 4: Evaluation of Reconstruction
Compute reconstruction errors ε_1, ε_2 for respective groups
Compute complementary reconstruction errors $\hat{\varepsilon_1}$, $\hat{\varepsilon_2}$
Derive reconstruction probabilities RP_1, RP_2
Step 5: Ink Mismatch Detection
if $RP_1 < \theta_2$ **and** $RP_2 < \theta_1$ **then**
    Infer ink mismatch present
else
    No ink mismatch detected
end if
Step 6: Clustering for Categorization
if mismatch detected **then**
    Compute total reconstruction errors $\mathcal{E}_1$, $\mathcal{E}_2$ for all ink pixels
    Apply K-means clustering on $[\mathcal{E}_1, \mathcal{E}_2]$ to categorize ink pixels
end if

Algorithm 1 Proposed Unsupervised Ink Mismatch Detection using Autoencoders

Autoencoders are particularly well-suited for this task due to their ability to learn compact and meaningful representations without requiring labeled data. Unlike more complex recent deep learning based models, which demand large labeled datasets and high computational overhead, autoencoders offer a lightweight, interpretable, and scalable solution, which is a key advantage for real-world spectral forensic document analysis. In order to validate our research hypothesis that spectral characteristics can be effectively captured and discriminated using autoencoders, we conducted a preliminary experiment with separate autoencoders trained on individual ink colours (blue and black). We then used these autoencoders to reconstruct spectra from both matching and non-matching ink samples. As shown in Fig. 3, reconstructions performed by the same-colour autoencoder yielded significantly lower reconstruction errors, whereas cross-colour reconstructions resulted in higher errors. This consistent difference supports the hypothesis that the autoencoder learns ink-specific spectral features.

This reconstruction error pattern is further exploited as an auxiliary signal for clustering. By analyzing reconstruction errors of multiple locally specialized autoencoders, ink samples can be grouped on the basis of their spectral compatibility. This enhances the robustness of the proposed method, allowing it to identify different ink types without relying on predefined models or labeled data.

4 Experimentation

This section discusses the experimental setup, including dataset descriptions, autoencoder configurations, experimental results, and discussions.

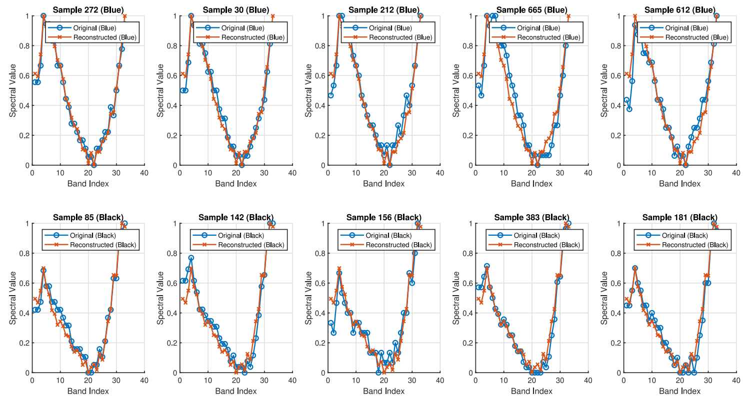

(a) Black and blue ink samples spectra reconstructed by their own (matched) autoencoders (i.e., blue AE on blue ink, black AE on black ink).

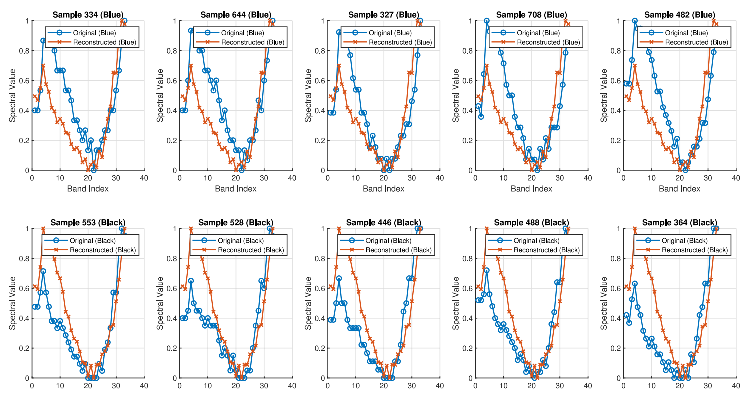

(b) Black and blue ink samples spectra reconstructed by the opposite (unmatched) autoencoders (i.e., blue AE on black ink, black AE on blue ink).

Fig. 3 Spectra reconstruction results using ink-specific autoencoders. Each plot displays the spectral signatures of randomly sampled blue and black ink pixels, where each line represents an individual pixel spectrum. **Top:** Reconstructions using their respective colour-trained autoencoders (i.e., matched inks). **Bottom:** Reconstructions using autoencoders trained on the opposite ink colour (i.e., unmatched inks). The x-axis represents the band index, and the y-axis indicates the spectral value. Higher reconstruction errors in the unmatched setting validate that autoencoders effectively learn ink-specific spectral features

4.1 Experimental setup

The University of Western Australia (UWA) Writing Ink Hyperspectral Image (WIHSI) dataset [20] is used to evaluate the proposed ink mismatch detection method. This dataset consists of seventy real handwritten notes, each one containing the phrase "*The quick brown fox jumps over the lazy dog*". Each image, captured by a hyperspectral sensor, is of size 81×627 pixels, with spectral bands ranging from 400 nm to 720 nm at intervals of 10 nm, yielding images with 33 spectral bands in the visible range. Seven people wrote the notes by hand on white paper with ten different pens, five blue and five black. This variety of pen types and ink colors makes it easier to do a full analysis of ink properties. Figure 4 shows a sample document image that was scanned across different spectral bands. The plotted curves correspond to approximately 50 foreground pixels of blue-coloured text. Additionally, Fig. 5 illustrates the corresponding spectra of foreground pixels within the scanned wavelength range of a representative hyperspectral document

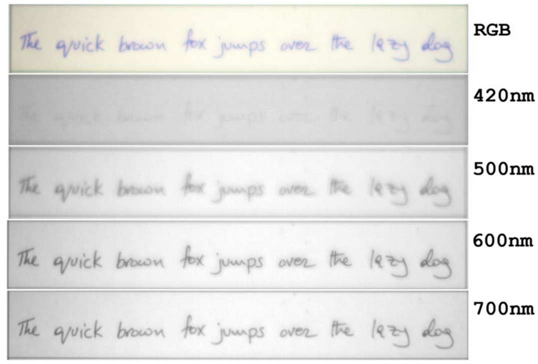

Fig. 4 A sample handwritten document of WIHSI dataset captured in RGB and some selected bands

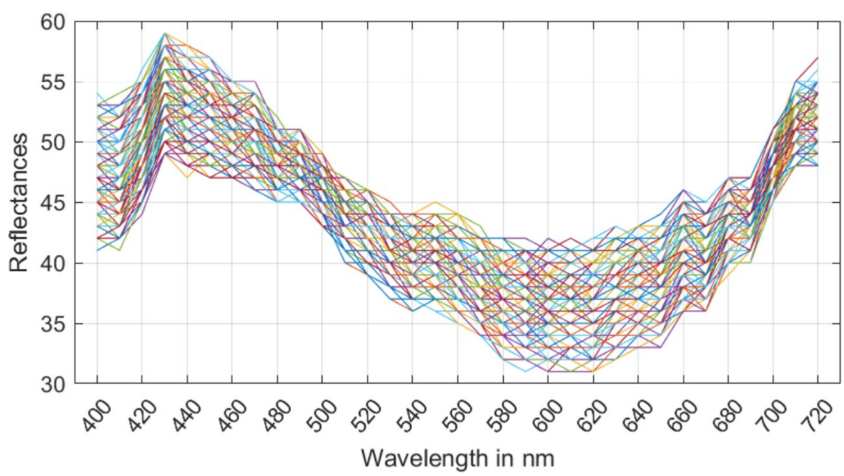

Fig. 5 Spectral signatures of foreground pixels extracted from a representative hyperspectral document image in the scanned wavelength range, corresponding to blue-coloured ink

image comprising 33 spectral bands. The plotted curves correspond to approximately 50 foreground pixels of blue-coloured text.

In order to simulate real-world scenarios of document manipulation involving different inks, multi-ink handwritten notes were prepared by merging segments of two randomly selected handwritten notes of the same ink colour, as detailed in [5, 20]. This procedure focuses solely on notes with the same ink colour, as it is uncommon for forgers to use different ink colours during manipulation, as such discrepancies can be readily detected visually. Moreover, this merging strategy was adopted to generate spectrally similar yet compositionally distinct ink regions, thereby enabling the proposed method to be evaluated under subtle and challenging conditions where discrimination cannot rely on apparent colour differences but rather on fine spectral variations. The preparation of multi-ink notes is depicted in Fig. 6, illustrating the merging process.

Hereafter, the symbol $\mathbf{C_{ij}}$ is used to denote a multi-ink sample representing a combination of proportions of different inks $\mathbf{i}$ and $\mathbf{j}$, resulting in ten distinct combinations per colour per subject.

The autoencoders employed in this work consist of an encoder and a decoder, each comprising multiple layers. The encoder takes the 33 spectral feature vectors from the segmented hyperspectral data as input. It then passes through two fully connected

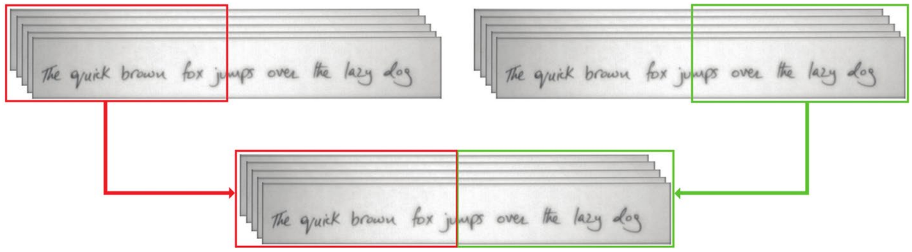

Fig. 6 Illustration for preparation of multi-ink handwritten notes

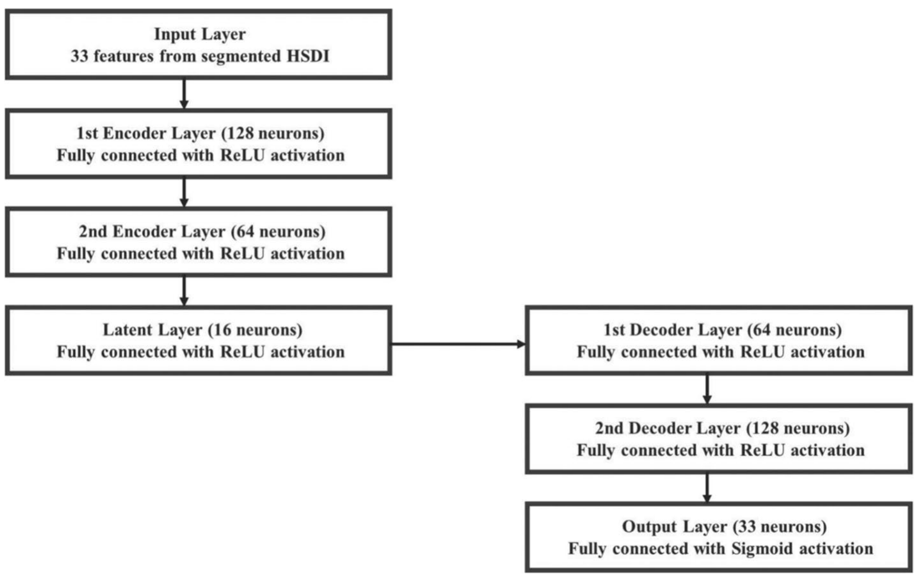

Fig. 7 Architecture of the autoencoder utilized in the proposed method

hidden layers: the first with 128 neurons and the second with 64 neurons, both using Rectified Linear Unit (ReLU) activation. The next layer of the encoder, the latent layer, is a fully connected layer with 16 neurons, also using ReLU activation. The decoder mirrors this structure, starting with a fully connected layer of 64 neurons, followed by a second hidden layer of 128 neurons with ReLU activation. The output layer, which has 33 neurons, uses the sigmoid activation function to reconstruct the input data. To prevent overfitting, dropout regularization with a rate of 0.1 is applied to the hidden layers. The model uses Xavier initialization [47] for stable training, and the loss function is Mean Squared Error (MSE), which is appropriate for continuous data reconstruction. The architecture of the autoencoders is depicted in Fig. 7.

The values of hyperparameters for the autoencoder are selected through grid search technique, where a range of hyperparameter combinations are exhaustively evaluated to identify the optimal configuration that maximizes detection accuracy. The selection of threshold values and window size is based on empirical analysis of the distribution of reconstruction errors obtained during training.

4.2 Experimental results

To evaluate the performance of the proposed method, multi-ink handwritten notes with equal ink proportions are prepared using the methodology discussed in the preceding section. The prepared document is partitioned spatially into two distinct segments using a sliding window mechanism traversing the document. Each segment is processed independently using two identical deep autoencoders. The entire procedure is repeated ten times, varying the position of the sliding window, and the final accuracy is computed as the average of all iterations.

First, the efficacy of the proposed method in detecting ink mismatches within the document is assessed. For the purpose, a total of 2,450[1] possible ink combination samples are considered; 2,380 samples are formed by combining portions belonging to different inks, and the remaining 70 samples are formed by combining portions from the same ink. Table 1 presents the performance evaluation of the proposed method for ink mismatch detection, depicting a confusion matrix alongside additional performance metrics such as precision, recall, accuracy, and the F_1 score. The performance metrics used for evaluation are defined as follows:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{3}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{4}$$

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{5}$$

$$F_1 \text{ Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{6}$$

where *TP*, *TN*, *FP*, and *FN* stand for the number of true positives, true negatives, false positives, and false negatives, respectively. The F_1 Score gives a balanced evaluation between precision and recall, while precision measures the proportion of correctly detected mismatch pixels among all predicted mismatches, and recall measures the proportion of actual mismatches correctly identified.

The performance parameters in Table 1 show that the proposed method works effectively for finding ink mismatches. The method shows robust capabilities to identify the difference between documents with and without ink mismatches, with high precision, recall, accuracy, and a F_1 score over 0.99. This comprehensive analysis shows that the proposed method for forensic document analysis is both reliable and accurate.

Next, the ink isolation accuracy of documents identified with ink mismatches is analyzed. To simplify the analysis and reduce experimental intricacies, the emphasis is placed exclusively on intra-subject multi-ink handwritten notes. We use the ink

Table 1 Performance of the proposed method for ink mismatch detection

	Ink mismatch (True)	**No ink mismatch (True)**
Ink Mismatch (Predicted)	2364	11
No Ink Mismatch (Predicted)	14	61
Precision	0.9954	
Recall	0.9941	
Accuracy	0.9898	
F_1 Score	0.9947	

[1] Allowing for repetitions and encompassing both intra-subject and inter-subject combinations $(25 \times 7 \times 7 \times 2 = 2,450)$

Table 2 Ink mismatch detection accuracies of the proposed method for different ink combinations

Blue ink combinations									
C_{12}	C_{13}	C_{14}	C_{15}	C_{23}	C_{24}	C_{25}	C_{34}	C_{35}	C_{45}
96.86	99.51	99.83	92.31	88.92	87.25	97.62	91.53	95.11	99.40
Mean ink mismatch detection accuracy for blue ink: 94.83									
Black ink combinations									
C_{12}	C_{13}	C_{14}	C_{15}	C_{23}	C_{24}	C_{25}	C_{34}	C_{35}	C_{45}
99.78	98.12	99.74	98.79	86.64	84.38	87.76	73.07	84.74	82.97
Mean ink mismatch detection accuracy for black ink: 89.59									

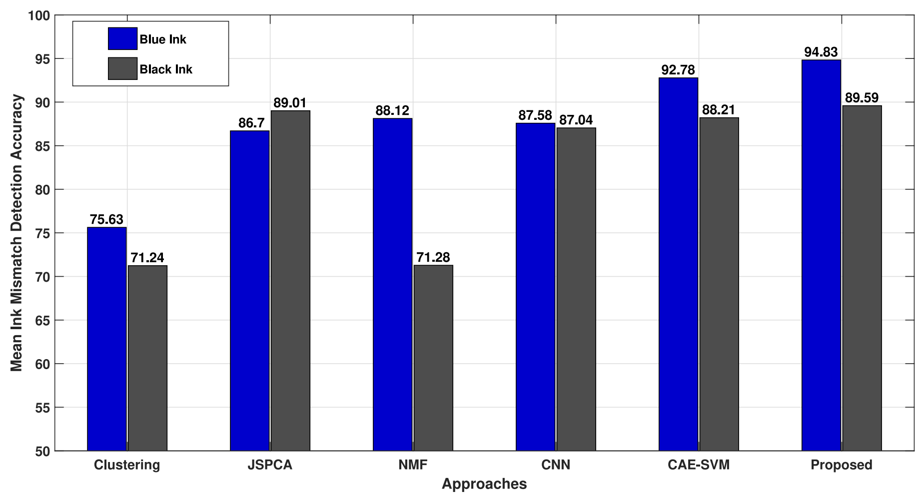

Fig. 8 Mean accuracy comparison of proposed approach with other state-of-the-art approaches: Clustering [20], Joint Sparse PCA [5], CNN [25], NMF [29], CAE-SVM [28], and proposed method

Table 3 Mean ink mismatch detection accuracy (%) for Blue and Black inks across different methods

Method	**Blue ink**	**Black ink**
Clustering [20]	75.63	71.24
Joint Sparse PCA [5]	86.70	89.01
NMF [29]	88.12	71.28
CNN [25]	87.58	87.04
CAE-SVM [28]	92.78	88.21
Proposed	94.83	89.59

differentiation accuracy metric from [5, 20] to see how well the proposed ink mismatch detection method works. To find the final accuracy for each ink combination, we take the average of the accuracies for all subjects. Table 2 summarizes the ink mismatch detection accuracies of the proposed method for different ink combinations for both blue and black ink. The comparison with results from state-of-the-art ink mismatch detection methods in [5, 20, 25, 28, 29] is illustrated in Fig. 8, with the corresponding mean accuracies summarized in Table 3.

The experimental results reveal that the proposed method consistently outperforms other state-of-the-art methods across various ink combinations. With accuracies reaching up to 94.83% for blue ink and 89.59% for black ink combinations, the proposed method demonstrates superior capability in accurately identifying ink mismatches. While some ink combinations exhibit a slight decline in performance, this can largely be attributed to the limited spectral information available in the dataset, which is restricted to the visible wavelength range and therefore lacks other band ranges like the near-infrared (NIR) bands that could capture finer spectral variations between similar inks. The consistently higher detection accuracy observed for blue inks (Table 2) compared to black inks can also be explained by several factors. Blue dyes typically have more distinctive spectral characteristics within the visible range, resulting in improved separability in spectral space. In contrast, black inks, often carbon-based, tend to produce near-flat spectral responses, causing reduced discriminability among different black samples. Extending the spectral coverage of the dataset by incorporating bands beyond the visible spectrum, for example, NIR bands, or employing colour-specific band selection may further improve the performance on black inks. Visual representations of the results obtained from the proposed method for isolating different inks in some randomly selected images are illustrated in Fig. 9. In these examples, the TRUE INK MAP images refer to the ground truth images, where blue and red colours are used solely for visualization purposes, representing different pens of the same ink colour. The RESULT images represent the output generated by the proposed method. The variations in colour represent distinct detected clusters corresponding to different pens. From the figure, we can conclude that even when the inks are visually identical, the proposed method can effectively identify spectral differences.

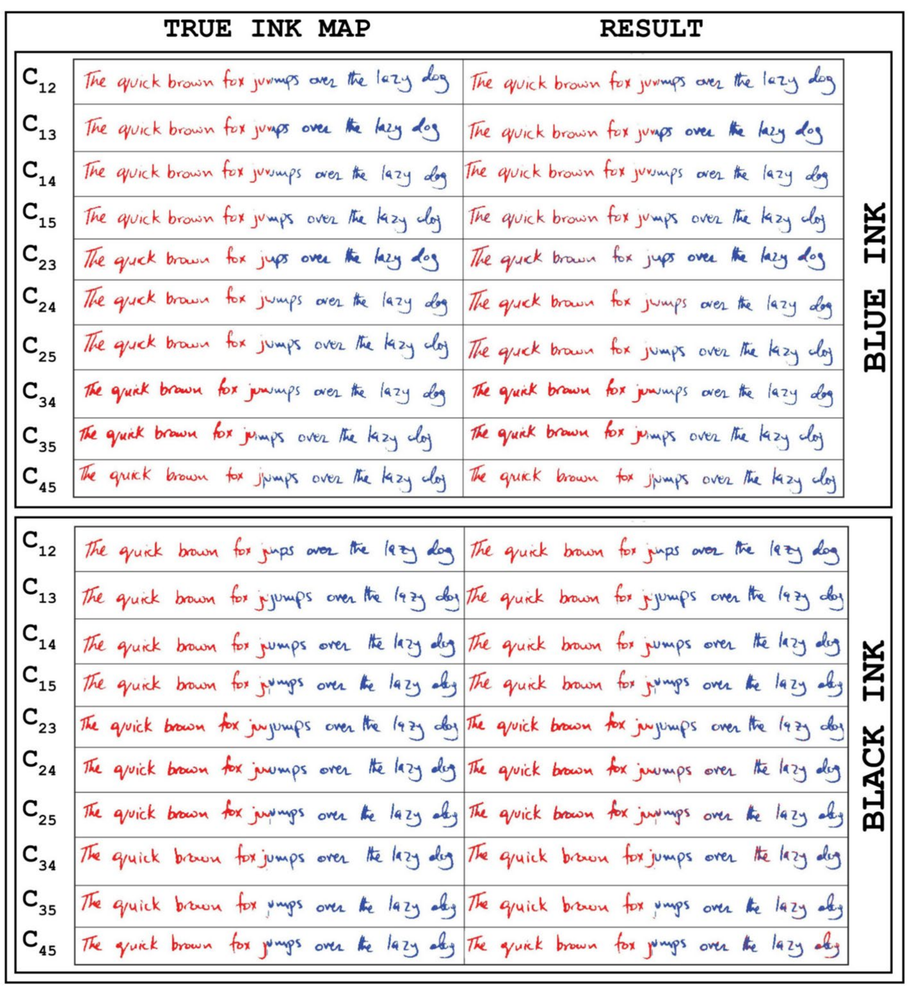

Fig. 9 Ink isolation results on some example images

In addition, two recent studies [30, 31] utilize Convolutional Autoencoder (CAE) features combined with Logistic Regression (LR) for ink mismatch detection and report high accuracies, 99.05% for blue inks and 92.74% for black inks in [30], and slightly improved accuracies in [31] using spectral-spatial features. However, these LR based approaches, being supervised methods, use labeled data and prior knowledge of ink classes, which limits their practical applicability in real-world forensic scenarios where inks and their types are often unknown. In contrast, the proposed unsupervised method in this study makes use of intrinsic spectral features to detect ink mismatches without the need for any labeled data. This allows robust detection of different ink types while maintaining interpretability and adaptability, thereby highlighting the practical advantages of the proposed method over other supervised approaches.

5 Ablation study: impact of parameter selection

To assess the robustness of the proposed ink mismatch detection method, we do an ablation study to look at how (i) the proportion of different inks, (ii) threshold values, and (iii) spatial segmentation parameters affect the performance of the proposed method.

5.1 Impact of disproportionate ink ratios

This subsection discusses how disproportionate ink ratios affect the accuracy of ink isolation accuracy in documents with mismatched ink. For the analysis, we examine multi-ink handwritten notes with four ink combination ratios (1:1, 1:2, 1:3, and 1:4), while varying the number of autoencoders employed for analysis from two to five. Figure 10

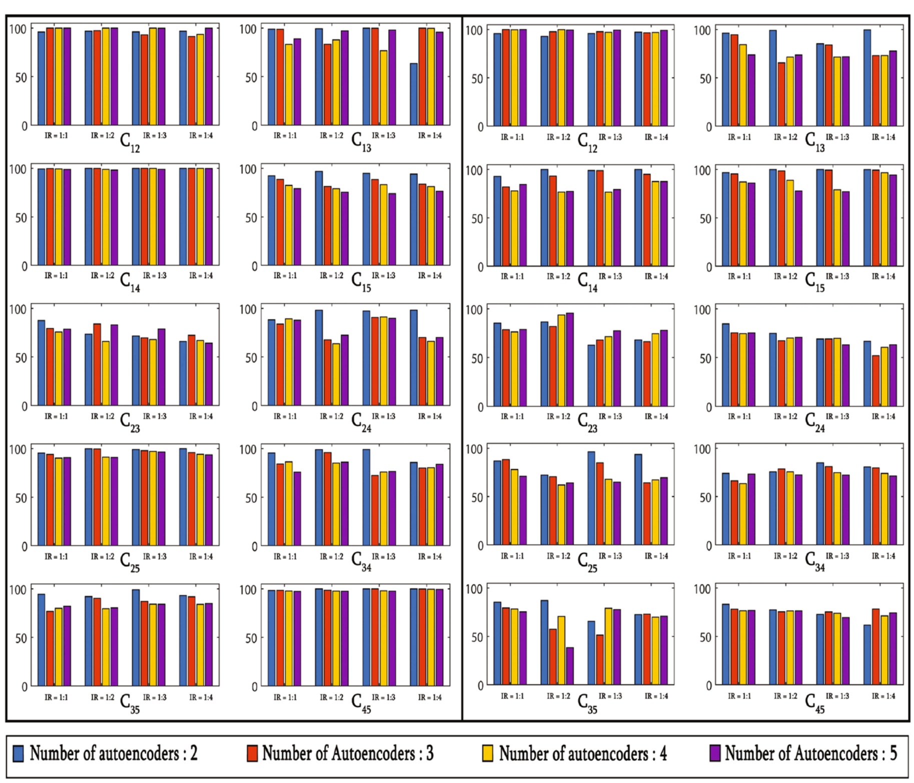

Fig. 10 Variation on multi-ink isolation accuracies with different ink ratios and different number of autoencoders for blue and black ink

shows how accurate the ink mismatch detection is for different ratios of ink combinations and different numbers of autoencoders.

The results show that using only two autoencoders is still effective in most cases. In some cases, accuracy goes down as the number of autoencoders goes up. This decrease in accuracy is likely due to under-complete training caused by insufficient spatially partitioned training data for the autoencoders. So, using two autoencoders gives the optimal results in our experimental setup for all tested ink ratios.

5.2 Automated parameter selection for threshold values

It is important to choose the right values for θ to find ink mismatches accurately. While identifying ink mismatches, the value of θ determines the sensitivity of the reconstruction error. Instead of relying on empirically observed values, we use the following three complementary methods to systematically explore and optimize the θ value:

- **Grid Search:** We look at a set range of θ values. The accuracy of ink mismatch detection is calculated for each value. This exhaustive strategy helps to visualize the performance trends and find the optimal region for θ.
- **Bayesian Optimization:** In order to improve efficiency, we use a probabilistic function to model how θ and accuracy are related. Bayesian optimization explores the parameter space iteratively by focusing on high-performing regions with fewer evaluations. This method is useful for finding values of θ that are near-optimal while keeping the cost of computation low.
- **Cross-Validation:** In order to ensure generalizability, k-fold cross-validation is applied across various document images and ink combinations. This step makes sure that the θ values chosen using the methods above work well on a wide range of data subsets.

Table 4 presents the detection accuracies for various values of θ, as determined through the above optimization strategies. Higher values typically improve precision by making the mismatch criteria stricter. Lower values, on the other hand, lead to more false positives. Through systematic parameter optimization and validation, we found that $\theta = 0.95$ for blue ink and $\theta = 0.90$ for black ink were the best trade-offs between sensitivity and robustness in detection.

5.3 Impact of window size and sliding step

The spatial segmentation parameters, which are the window size and sliding step, decide how the document is partitioned for analysis by the autoencoders. Choosing the right values for these parameters affects the granularity of feature extraction and mismatch localization, which has a big impact on how accurate ink mismatch detection is. We looked at a number of different setups to analyze the impact. The window sizes were set

Table 4 Impact of threshold values on ink mismatch detection accuracy

Threshold value (θ)	Blue ink accuracy	Black ink accuracy
0.60	75.63	71.24
0.70	86.70	89.01
0.80	88.12	71.28
0.85	87.58	87.04
0.90	92.78	88.21
0.95	94.79	89.13

Table 5 Impact of sliding window size on ink mismatch detection accuracy

Window size	Sliding step	Black ink accuracy
20%	5%	93.7%
30%	10%	96.2%
40%	15%	92.4%
50%	20%	89.6%

to 20%, 30%, 40%, and 50% of the document width, and the sliding steps were set to 5%, 10%, 15%, and 20%.

Table 5 shows a summary of the results. The analysis shows that using smaller window sizes and finer sliding steps makes it easier to identify ink because it allows for more localized reconstruction analysis. It should be noted that they also introduce additional computational overhead. On the other hand, larger windows and coarser steps lower processing costs, but they miss small differences in the spectra of different inks, which can make performance worse. We used Bayesian optimization to identify the best balance between detection accuracy and computational efficiency. Bayesian optimization lets us identify the best combinations of parameters without having to search the whole space. The chosen settings of a 30% window size and a 10% sliding step gave the best balance between precision and efficiency.

These results show the robustness of the proposed approach across varying parameter settings and validate the utility of systematic parameter tuning techniques over empirical selection. This also lays the foundation for future expansions, where adaptive parameter selection mechanisms based on document complexity or ink distribution could dynamically change key parameters during deployment, making the system more flexible and scalable.

Effect of Paper Material All samples in the dataset were written on the same type of standard white paper. As the proposed method uses Sauvola's local thresholding for ink-pixel segmentation prior to spectral analysis, the effect of paper material on detection accuracy is inherently minimized. Since only the segmented ink pixels are utilized for further processing, variations in paper background should not contribute significantly to the overall performance.

6 Conclusion and future directions

This work presents a novel deep autoencoder-based approach for ink mismatch detection in HSDIs using a completely unsupervised framework. The main strength of the proposed approach is its fully unsupervised design that works without any labeled data or prior spectral information that makes it particularly well-suited for real-world forensic document analysis, where such information is often unavailable. By utilizing the latent feature learning ability of autoencoders, the method is able to detect subtle spectral differences between ink pixels without any prior spectral information or labeled training data. Experimental results demonstrate the superiority of the proposed method over existing ones across various ink combinations. Additionally, ablation studies validate its robustness on different ink ratios, threshold values, and segmentation strategies. Nevertheless, there are some factors that can affect the efficiency of the proposed method. For instance, keeping the lighting condition stable during image capturing is recommended because variations in light intensity across the document can alter the recorded spectral characteristics of inks. The proposed method is also able to detect multiple ink types, but it may still be hard to tell the difference if the document consist numerous similar inks with complex ink compositions. The empirical selection of windows during experimentation works well, but it could be improved even more. Automating the selection of model parameters through systematic optimization will make the method easier and more reliable. Future efforts will focus on addressing these areas to enhance overall adaptability. Future extensions may also consider hybrid methodologies, including self-supervised contrastive learning, transformer-based models, or diffusion-based generative frameworks to improve its adaptability and scalability. These could improve latent feature discrimination and generalization, especially in complex or resource-constrained forensic settings. Expanding the spectral range beyond the visible spectrum (e.g., into NIR or UV) may also enhance detection accuracy. Furthermore, future investigations will include robustness analysis of the proposed framework against common post-processing operations typically encountered in Online Social Network (OSN) environments, such as compression, resizing, and filtering, to assess its reliability under real-world image sharing conditions. Finally, broader validation on additional datasets, when publicly available, will be essential to establish the generalizability of the proposed framework.

Acknowledgements
The authors gratefully acknowledge Manipal University Jaipur, Banaras Hindu University, and Doctor Harisingh Gour Vishwavidyalaya, Sagar (M.P.) for their support, infrastructure, and conducive research ecosystem, all of which played a crucial role in the successful completion of this work.

Author contributions
Pangambam Sendash Singh: Conceptualization, original draft preparation. Subbiah Karthikeyan: Overall supervision of the study. Govind Murari Upadhayay: Data collection and analysis. Sounak Sadhukhan: Contributed to literature review and manuscript editing. Pramod Kumar Soni: Critical revision of the manuscript, administrative support. Timothy Malche: Interpretation of data and results.

Funding
Open access funding provided by Manipal University Jaipur.

Data availability
The datasets analysed during the current study are available in the University of Western Australia repository and can also be accessed from the dataset author's homepage https://sites.google.com/site/zohaibnet/Home/databases.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publcation
Not applicable.

Competing interests
The authors declare that they have no conflict of interest.

Received: 20 July 2025 / Accepted: 27 November 2025
Published online: 28 December 2025

References

1. Govender M, Chetty K, Bulcock H. A review of hyperspectral remote sensing and its application in vegetation and water resource studies. Water SA. 2009;33(2):145.
2. Selci S. The future of hyperspectral imaging. J Imag. 2019;5(11):84.
3. Khan MJ, Khan HS, Yousaf A, Khurshid K, Abbas A. Modern trends in hyperspectral image analysis: a review. IEEE Access. 2018;6:14118–29.
4. Edelman GJ, Gaston E, Leeuwen TG, Cullen PJ, Aalders MCG. Hyperspectral imaging for non-contact analysis of forensic traces. Forensic Sci Int. 2012;223(1–3):28–39.
5. Khan Z, Shafait F, Mian A. Automatic ink mismatch detection for forensic document analysis. Pattern Recogn. 2015;48(11):3615–26.
6. Dégardin K, Desponds A, Roggo Y. Protein-based medicines analysis by Raman spectroscopy for the detection of counterfeits. Forensic Sci Int. 2017;278:313–25.

7. Oravec M, Beganović A, Gál L, Čeppan M, Huck CW. Forensic classification of black inkjet prints using Fourier transform near-infrared spectroscopy and Linear Discriminant Analysis. Forensic Sci Int. 2019;299:128–34.
8. Miskelly GM, Wagner JH. Using spectral information in forensic imaging. Forensic Sci Int. 2005;155(2–3):112–8.
9. Bentolila A, Hartman I, Levin-Elad M. Blood or not blood that is the question. A non-destructive method for the detection of blood-contaminated fingermarks. Forensic Sci Int. 2017;278:374–8.
10. Bergmann T, Heinke F, Labudde D. Towards substrate-independent age estimation of blood stains based on dimensionality reduction and k-nearest neighbor classification of absorbance spectroscopic data. Forensic Sci Int. 2017;278:1–8.
11. Aginsky VN. Forensic examination of slightly soluble ink pigments using thin-layer chromatography. J Forensic Sci. 1993;38(5):13516.
12. Joo Kim S, Deng F, Brown MS. Visual enhancement of old documents with hyperspectral imaging. Pattern Recogn. 2011;44(7):1461–9.
13. Khan MJ. Deep learning for automated forgery detection in hyperspectral document images. J Electron Imaging. 2018;27(05):1.
14. Khan Z, Shafait F, Mian A. Hyperspectral document imaging: challenges and perspectives. In: International workshop on camera-based document analysis and recognition, 2013; pp. 150–63. Springer.
15. Qureshi R, Uzair M, Khurshid K, Yan H. Hyperspectral document image processing: Applications, challenges and future prospects. Pattern Recogn. 2019;90:12–22.
16. Sukhija R, Kumar M, Jindal M. Document forgery detection: a comprehensive review. Int J Data Sci Analyt. 2025;20(5):4385–407.
17. Brauns EB, Dyer RB. Fourier transform hyperspectral visible imaging and the nondestructive analysis of potentially fraudulent documents. Appl Spectrosc. 2006;60(8):833–40.
18. Silva CS, Pimentel MF, Honorato RS, Pasquini C, Montalbán JMP, Ferrer A. Near infrared hyperspectral imaging for forensic analysis of document forgery. Analyst. 2014;139(20):5176–84.
19. Morales A, Ferrer MA, Diaz-Cabrera M, Carmona C, Thomas GL: The use of hyperspectral analysis for ink identification in handwritten documents. In: 2014 international Carnahan conference on security technology (ICCST), IEEE, Rome, Italy. 2014; pp. 1–5.
20. Khan Z, Shafait F, Mian A. Hyperspectral imaging for ink mismatch detection. In: 2013 12th international conference on document analysis and recognition. IEEE, Washington, DC, USA, 2013; pp. 877–81.
21. Khan MJ, Yousaf A, Khurshid K, Abbas A, Shafait F. Automated forgery detection in multispectral document images using fuzzy clustering. In: 2018 13th IAPR international workshop on document analysis systems (DAS), 2018; pp. 393–8. IEEE.
22. Luo, Z., Shafait, F., Mian, A.: Localized forgery detection in hyperspectral document images. In: 2015 13th international conference on document analysis and recognition (ICDAR), 2015; pp. 496–500. IEEE, Nancy, France
23. Abbas A, Khurshid K, Shafait F. Towards automated ink mismatch detection in hyperspectral document images. In: 2017 14th IAPR international conference on document analysis and recognition (ICDAR), vol. 1, 2017; pp. 1229–36. IEEE, Kyoto, Japan.
24. Roy P, Bag S. Ink analysis based forensic investigation of handwritten legal documents. Multimed Tools Appl. 2022;81(16):23007–47.
25. Islam AU, Khan MJ, Khurshid K, Shafait F. Hyperspectral image analysis for writer identification using deep learning. In: 2019 digital image computing: techniques and applications (DICTA). IEEE, Perth, Western Australia, 2019; pp. 1–7.
26. Khan MJ, Khurshid K, Shafait F. A spatio-spectral hybrid convolutional architecture for hyperspectral document authentication. In: Proceedings of the international conference on document analysis and recognition, ICDAR, 2019; pp. 1097–102.
27. Jaiswal G, Sharma A, Kumar Yadav S. DFD-SS: document forgery detection using spectral spatial features for hyperspectral images. J Vis Commun Image Represent. 2022;89(November):103690.
28. Tyagi P, Agarwal K, Jaiswal G, Sharma A, Rani R. Forged document detection and writer identification through unsupervised deep learning approach. Multimed Tools Appl. 2024;83(6):18459–78.
29. Rahiche A, Cheriet M. Forgery detection in hyperspectral document images using graph orthogonal nonnegative matrix factorization. In: 2020 IEEE/CVF conference on computer vision and pattern recognition workshops (CVPRW), vol. 2020. IEEE, Washington, USA, 2020; pp. 2823–31.
30. Jaiswal G, Sharma A, Yadav SK. Deep feature extraction for document forgery detection with convolutional autoencoders. Comput Electr Eng. 2022;99:107770.
31. Jaiswal G, Sharma A, Yadav SK. Dfd-ss: document forgery detection using spectral-spatial features for hyperspectral images. J Vis Commun Image Represent. 2022;89:103690.
32. Nasir FA, Liaquat S, Khurshid K, Mahyuddin NM. A hyperspectral unmixing approach for ink mismatch detection in unbalanced clusters. J Inform Intel. 2024;2(2):177–90.
33. Nasir FA, Liaquat S, Khurshid K, Mahyuddin NM. A hyperspectral unmixing approach for ink mismatch detection in unbalanced clusters. J Inform Intell. 2024;2(2):177–90.
34. Petrík M, Mataš E, Sabo M, Ries M, Matejčík Š. Fast detection and classification of ink by ion mobility spectrometry and artificial intelligence. IEEE Access. 2025;13:33379–86.
35. Sun K, Cao G, Zhao Q, Zhang J. Differential abnormality-based tampering detection in digital document images. In: 2019 IEEE/ACIS 18th international conference on computer and information science (ICIS), 2019; pp. 145–9.
36. Zheng S, Chen C, Chen S, Yao T, Ding S, Li B, Huang J. Generalized document tampering localization via color and semantic disentanglement. IEEE Trans Circuits Syst Video Technol. 2025. https://doi.org/10.1109/TCSVT.2025.3597602
37. Lou Z, Cao G, Guo K, Yu L, Weng S. Exploring multi-view pixel contrast for general and robust image forgery localization. IEEE Trans Inf Forensics Secur. 2025;20:2329–41.
38. Zhu H, Cao G, Huang X. Progressive feedback-enhanced transformer for image forgery localization. arXiv:2311.08910, 2025.
39. Ballard DH. Modular learning in neural networks. In: Conference on artificial intelligence, 1987; pp. 279–84.
40. Chen Y, Lin Z, Zhao X, Wang G, Gu Y. Deep learning-based classification of hyperspectral data. IEEE J Sel Top Appl Earth Observ Remote Sens. 2014;7(6):2094–107.
41. Goodfellow I, Bengio Y, Courville A. Deep learning. Cambridge, MA, USA: MIT Press; 2016.
42. Liu X, Gherbi A, Wei Z, Li W, Cheriet M. Multispectral image reconstruction from color images using enhanced variational autoencoder and generative adversarial network. IEEE Access. 2021;9:1666–79.

43. Jiang J, Ren H, Zhang M. A convolutional autoencoder method for simultaneous seismic data reconstruction and denoising. IEEE Geosci Remote Sens Lett. 2022;19:1.
44. Sauvola J, Pietikäinen M. Adaptive document image binarization. Pattern Recogn. 2000;33(2):225–36.
45. Otsu N. A threshold selection method from gray-level histograms. IEEE Trans Syst Man Cybern. 1979;9(1):62–6.
46. Guo Y, Liao W, Wang Q, Yu L, Ji T, Li P. Multidimensional time series anomaly detection: a gru-based Gaussian mixture variational autoencoder approach. Proc Mach Learn Res. 2018;95(2001):97–112.
47. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the thirteenth international conference on artificial intelligence and statistics, 2010; pp. 249–56. JMLR Workshop and Conference Proceedings.

Publisher's Note