JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

MACHINE LEARNING FOR HANDWRITTEN DIGIT RECOGNITION: ENHANCING J2 CRITERION FEATURE COMPRESSION WITH MULTI-DESCRIPTOR FUSION

Shahzeb Jadoon[1], Muqaddas Yousaf [2], Shanza Mehboob [3], Muhammad Asif [4], Sohair Sultan Fayyaz Jeelani [5], Hadiya Ali [6]

Affiliations

[1, 2, 3, 4, 5, 6] Department of Computer Science, Abbottabad University of Science and Technology (AUST), Pakistan

Corresponding Author's Email

[5] sohair.sultan2266@gmail.com

ABSTRACT

This paper extends a prior study on J2 criterion feature compression for handwritten digit recognition by investigating three enhancement directions identified as future work: spatial pyramid extension of Local Binary Patterns, the effect of training set size on compression quality, and multi-descriptor feature fusion incorporating Scale Invariant Feature Transform features. The baseline pipeline extracts hand-crafted features, compresses them into a nine-dimensional discriminant subspace using a from-scratch implementation of the J2 criterion, and classifies the result using a Support Vector Machine and Linear Discriminant Analysis. Spatial Pyramid Local Binary Pattern encoding raises the accuracy of texture-only recognition from 56.25 percent to 92.70 percent by preserving the spatial arrangement of local texture that a global histogram discards. A data scaling study across 5,000 to 60,000 training samples shows that recognition accuracy improves from 94.15 percent to 96.75 percent while the J2 criterion value remains approximately stable, confirming that the compression quality is largely data-volume independent for a fixed number of classes. A triple-descriptor fusion combining Spatial Pyramid Local Binary Pattern, Histogram of Oriented Gradients, and Scale Invariant Feature Transform Bag of Words features achieve the highest accuracy of 97.30 percent with the largest J2 value of 38.17, demonstrating that the three descriptors carry complementary discriminative information. The results validate the J2 compression framework as an effective and scalable dimensionality reduction approach across all tested feature configurations.

Keywords: Handwritten Digit Recognition, J2 Criterion, Spatial Pyramid, Local Binary Pattern, Histogram of Oriented Gradients, Sift, Feature Fusion, Support Vector Machine, Dimensionality Reduction.

I. INTRODUCTION

Handwritten digit recognition is one of the most popular problems in the field of computer vision and pattern recognition, as it is not only a useful tool for postal automation and bank cheque processing and for form digitisation but it is also a good testbed for the evaluation of feature extraction and classification techniques [1] [2]. Several decades of methodological advances have occurred on the MNIST dataset, which consists of 70,000 greyscale images of handwritten digits with a resolution of 28 x 28 pixels, ranging from classical feature engineering pipelines to modern deep learning approaches based on convolutional architectures [3]. Although deep learning approaches today can almost match humans on MNIST, there are still significant applications for classical pipelines based on hand-crafted features and interpretable dimensionality reduction in resource-limited settings, applications where model transparency is crucial, and as pedagogical tools to understand the relationship between feature quality and classifier performance [4] [5].

JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

In a previous work, a classical recognition pipeline was defined, which included the use of Local Binary Pattern and Histogram of Oriented Gradients feature extraction methods along with a Support Vector Machine and Linear Discriminant Analysis classifiers, and implemented the feature compression using the J2 criterion from scratch. That work showed that the J2 criterion, based on Fisher's discriminant analysis [6] [7] is effective in projecting high dimensional feature vectors into a compact nine dimensional discriminant subspace in such a way that the class-relevant structure is retained. The study found that Local Binary Pattern (LBP) features are inferior to Histogram of Oriented Gradients (HOG) features, with recognition accuracy of 56.25 percent versus 95.55 percent for digit recognition, and that the best accuracy of 95.95 percent was achieved using a concatenated LBP and HOG (LBPhog) descriptor. Four directions for further research were suggested in the previous work, and the present paper systematically addresses three of them.

The first drawback of the baseline Local Binary Pattern (LBP) descriptor is that it constructs a single global histogram of an image's texture information, thereby losing all spatial arrangement information [8]. Two totally different digits can have similar global texture histograms, because their local neighbourhood distributions are similar, which is why the baseline descriptor does not distinguish well. Lazebnik et al. [9] overcame this by dividing the image into increasingly finer spatial pyramids and computing a histogram for each cell, which provides a spatial representation of local texture features for scene classification. By applying this principle to Local Binary Patterns, the spatial information important for digit-class separation can significantly improve the accuracy of digit recognition [10].

The second direction is on the effect of training set size on the quality of J2 compression and classification accuracy. The number of training samples in the training partition was deliberately reduced to 10,000 in the baseline study to eliminate the impact of large-scale training samples on feature quality. However, the scaling behaviour of the J2 criterion value was not described by the number of training data. It is important to understand this relationship because the J2 criterion estimates scatter matrices from the training data and the stability of the estimate as a function of the number of samples dictates the reliability of the compression under conditions of limited training data during deployment [11].

The third direction is the introduction of a third complementary feature descriptor. To complement the dense texture and gradient features of the baseline, the Scale Invariant Feature Transform [12] is used to detect and describe local keypoints in scale- and rotation-invariant ways, and their aggregation is represented as a histogram of Bag-of-Words [13] [14]. Multi-descriptor fusion has consistently been shown to be effective at boosting recognition accuracy compared to a single descriptor, as it combines information from different levels of abstraction [15] [16].

This paper makes the following contributions. Then, a Spatial Pyramid Local Binary Pattern (SPLP) descriptor is proposed and demonstrated to achieve 92.70 percent recognition accuracy on texture-only images, thereby reducing the performance gap between texture-only and gradient-based descriptors to a minimum. Second, a systematic data-scaling study across four training set sizes characterises the relationships among training set size, the J2 criterion value, and recognition accuracy. Third, a fusion of spatial pyramid local binary pattern, histogram of oriented gradients, and scale-invariant feature transform bag-of-words features achieves the best reported accuracy of 97.30 per cent. The J2 compression stage, the Support Vector Machine, and the Linear Discriminant Analysis classifier remain the same throughout, as they were in the baseline study for a controlled and fair comparison. The remainder of this paper is organised as follows. The proposed methodology and each feature descriptor is described in Section 2. The result of the experiments and discussion are presented in section 3. The paper ends with the conclusions presented in section 4 and future work underway.

II. PROPOSED METHODOLOGY

The proposed recognition pipeline retains the three-stage structure of the baseline study: feature extraction, feature compression using the J2 criterion, and classification. The present work extends the feature extraction stage with two new descriptors, a Spatial Pyramid Local Binary Pattern and a Scale Invariant Feature Transform Bag of Words, and introduces a triple-descriptor fusion configuration. The compression and classification stages are carried over unchanged to enable a controlled comparison. Fig. 1 illustrates the overall pipeline architecture.

Fig. 1: OVERALL ARCHITECTURE Of THE PROPOSED RECOGNITION PIPELINE. THE FEATURE EXTRACTION STAGE IS EXTENDED WITH SPATIAL PYRAMID LBP AND SIFT BAG OF WORDS DESCRIPTORS, WHILE THE J2 COMPRESSION AND CLASSIFICATION STAGES ARE CARRIED OVER UNCHANGED FROM THE BASELINE STUDY.

Input Image 28×28	**Feature** Extraction (SP-LBP / HOG / SIFT-BoW)	**Standardise** (Z-score)	**J2 Compression** 342→9 dims (from scratch)	**Classification** (SVM / LDA)	**Digit** Prediction

A. Baseline Feature Descriptors

The baseline Local Binary Pattern (LBP) descriptor [8] encodes local texture by comparing each pixel with its circular neighbourhood of P equally spaced sampling points at radius r. For a central pixel with intensity pc, the LBP code is computed as

$$LBP(P, r) = \Sigma\, s(p_i - pc) \times 2^i, \quad s(x) = 1 \text{ if } x \geq 0, \text{ else } 0 \qquad (1)$$

where the summation runs over the P neighbours. The uniform variant is used, producing an 18-dimensional normalised histogram at P = 16 and r = 2. The Histogram of Oriented Gradients (HOG) descriptor [15] captures the local distribution of gradient orientations over a spatial grid of cells, producing a 324-dimensional feature vector for a 28 by 28 image with 7 by 7 pixel cells, 9 orientation bins, and 2 by 2 block normalisation.

B. Spatial Pyramid Local Binary Pattern

The principal weakness of the global LBP histogram is that it discards all spatial information by aggregating texture codes across the entire image. The proposed Spatial Pyramid Local Binary Pattern (SP-LBP) descriptor addresses this by adopting the spatial pyramid matching principle [9]. The LBP code map is first computed for the whole image. The map is then partitioned into a set of grids at multiple pyramid levels, and a separate normalised LBP histogram is computed for each cell. The per-cell histograms are concatenated to form the final descriptor.

Let L denote the set of pyramid levels. At level ℓ, the image is divided into ℓ × ℓ non-overlapping cells. The total number of cells across all levels is

$$N\_cells = \Sigma\, \ell^2 \text{ for } \ell \in L \qquad (2)$$

and the final feature dimension is (P + 2) × N_cells. With pyramid levels L = {1, 2, 3}, the number of cells is 1 + 4 + 9 = 14, and the resulting descriptor has dimension 18 × 14 = 252. This preserves the spatial arrangement of local texture patterns while retaining the illumination robustness of the underlying LBP operator. Algorithm 1 summarises the extraction procedure.

Algorithm 1: Spatial Pyramid LBP Feature Extraction

```
Input: image I (28×28), points P, radius r, levels L = {1, 2, 3}
Output: feature vector f
M ← uniform_LBP(I, P, r)          // compute LBP code map
f ← empty list
for each level ℓ in L do
    (ch, cw) ← (H/ℓ, W/ℓ)          // cell height, width
    for r = 0 to ℓ−1 do
        for c = 0 to ℓ−1 do
            cell ← M[r·ch:(r+1)·ch, c·cw:(c+1)·cw]
            h ← histogram(cell, P+2 bins)
            h ← h / (sum(h) + ε)      // L1 normalise
            append h to f
return concatenate(f)             // dimension (P+2)×Σℓ²
```

JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

C. SIFT Bag of Words Descriptor

The Scale Invariant Feature Transform (SIFT) [12] detects local keypoints and describes each by a 128-dimensional gradient orientation histogram that is invariant to scale and rotation. Because the number of detected keypoints varies per image, the descriptors are aggregated into a fixed-length representation using the Bag of Words model [13], [14]. A visual vocabulary of K = 80 words is first learned by clustering the SIFT descriptors extracted from a subset of the training images using mini-batch K-means. Each image is then represented by a normalised histogram over the K visual words

$$h_k = (1 / N) \Sigma\, \delta(\text{assign}(d_i) = k), \quad k = 1, ..., K \qquad (3)$$

where N is the number of keypoints in the image, d_i is the i-th SIFT descriptor, and assign(•) returns the index of the nearest visual word. The resulting 80-dimensional histogram provides a keypoint-based representation that is complementary to the dense texture and gradient descriptors.

D. Multi-Descriptor Fusion

To exploit the complementary nature of the three descriptors, feature-level fusion is performed by concatenating the individual feature vectors prior to compression. Two fusion configurations are evaluated: a triple fusion of baseline LBP, HOG, and SIFT Bag of Words with dimension 18 + 324 + 80 = 422, and a full fusion of SP-LBP, HOG, and SIFT Bag of Words with dimension 252 + 324 + 80 = 656. In both cases the concatenated vector is standardised and then compressed by the J2 criterion to a nine-dimensional discriminant subspace.

E. J2 Criterion Feature Compression

The J2 criterion, carried over unchanged from the baseline study, is a supervised dimensionality reduction method rooted in Fisher's discriminant analysis [6], [7]. Given a labelled training set, two scatter matrices are computed: the within-class scatter matrix Sw and the between-class scatter matrix Sb:

$$S_{vv} = \Sigma c\ \Sigma_i{\in}c\ (x_i - \mu c)(x_i - \mu c)^T \qquad (4)$$

$$S^b = \Sigma c\ nc\ (\mu c - \mu)(\mu c - \mu)^T \qquad (5)$$

where μc is the mean of class c, μ is the global mean, and nc is the number of samples in class c. The J2 criterion value quantifies class separability as

$$J_2 = \text{trace}(S_{vv}^{-1}\ S^b) \qquad (6)$$

Maximising J2 leads to the generalised eigenvalue problem Sw^{-1} Sb w = λ w, whose top C − 1 = 9 eigenvectors form the projection matrix that maps each feature vector to a nine-dimensional discriminant representation. A principal component analysis pre-processing step retaining 98 percent of variance is applied to high-dimensional descriptors to ensure the within-class scatter matrix is well-conditioned before inversion.

F. Classification

Two classifiers are evaluated on the compressed features, both carried over unchanged from the baseline. The Support Vector Machine (SVM) [17] uses a radial basis function kernel with regularisation parameter C = 10 and a one-versus-rest multi-class strategy. Linear Discriminant Analysis (LDA) [11] provides a computationally efficient linear classifier that is well matched to the J2-compressed subspace, which is designed to maximise linear separability. Table I summarises the feature configurations and their dimensionalities.

Table 1: FEATURE CONFIGURATIONS AND THEIR DIMENSIONALITIES BEFORE AND AFTER J2 COMPRESSION

Feature Configuration	Type	Original Dim	J2 Dim
LBP	Texture (global)	18	9
SP-LBP	Texture (spatial)	252	9
HOG	Gradient	324	9
SIFT-BoW	Keypoint	80	9
LBP+HOG	Fusion (baseline best)	342	9
LBP+HOG+SIFT	Fusion (triple)	422	9
SP-LBP+HOG+SIFT	Fusion (full)	656	9

JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

Table 2: EXPERIMENTAL CONFIGURATION USED ACROSS ALL REPORTED EXPERIMENTS.

Parameter	Value
Test set size	2,000 samples
LBP points / radius	P = 16, r = 2
SP-LBP pyramid levels	L = {1, 2, 3} (14 cells)
HOG cells / orientations	7×7 px, 9 bins
SIFT vocabulary size	K = 80 visual words
J2 target dimension	9 (C − 1)
SVM kernel / C	RBF, C = 10
Random seed	42

III. RESULTS AND DISCUSSION

All experiments use a fixed test set of 2,000 samples drawn from MNIST, and the J2-compressed nine-dimensional feature space is used throughout to allow a direct and fair comparison. Results are organised into three experiments corresponding to the three enhancement directions.

A. Experiment 1 Spatial Pyramid LBP

Fig. 2 compares the recognition accuracy and J2 criterion value of the baseline LBP, the proposed SP-LBP, and the reference HOG descriptor. The baseline global LBP achieves only 56.25 percent accuracy with an SVM, reflecting the poor class separability of the 18-dimensional global histogram with a J2 value of 4.01. The proposed SP-LBP achieves accuracy of 92.70 percent with an SVM and 91.60 percent with LDA, while increasing the J2 value by more than threefold to 15.04. This confirms that preserving the spatial arrangement of local texture recovers the majority of the discriminative information lost by global aggregation, closing most of the gap with HOG's 95.55 percent accuracy.

Fig. 2: COMPARISON OF BASELINE LBP, SPATIAL PYRAMID LBP, AND HOG FEATURES. LEFT: RECOGNITION ACCURACY FOR SVM AND LDA CLASSIFIERS. RIGHT: J2 CRITERION VALUE BY FEATURE TYPE. SP-LBP RAISES ACCURACY FROM 56.25 PERCENT TO 92.70 PERCENT BY ENCODING SPATIAL TEXTURE ARRANGEMENT.

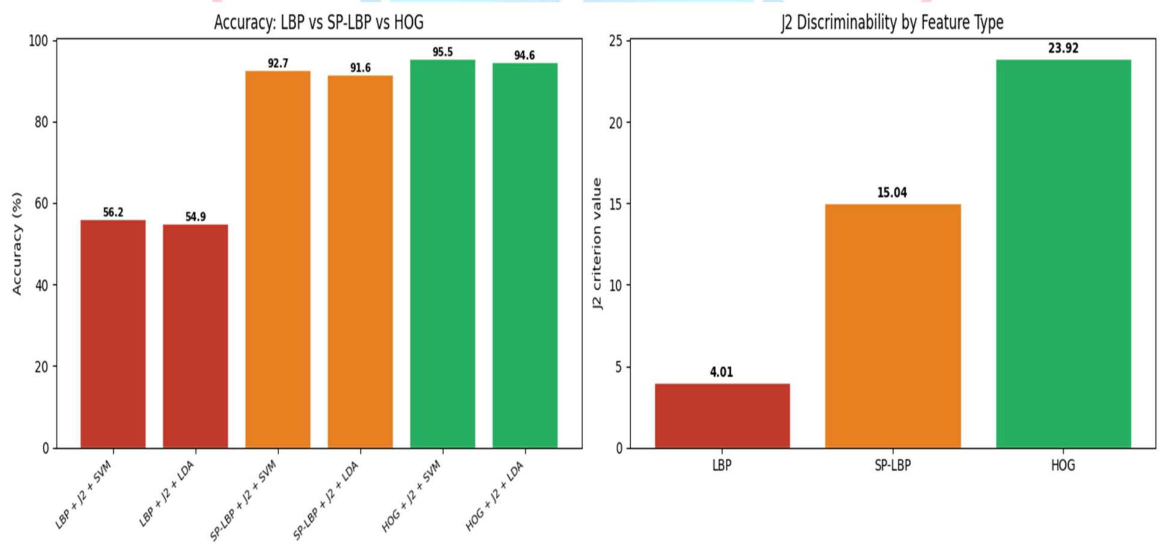

Table III reports the exact accuracy values. The 36.45 percentage-point improvement from LBP to SP-LBP is the largest gain observed in this study, demonstrating that the spatial pyramid extension is a highly effective, low-cost enhancement to the texture descriptor.

JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

Table 3: EXPERIMENT 1 RESULTS: ACCURACY OF BASELINE LBP, SP-LBP, AND HOG WITH J2 COMPRESSION.

Feature	J2 Value	SVM Acc. (%)	LDA Acc. (%)
LBP (baseline)	4.01	56.25	54.95
SP-LBP (proposed)	15.04	92.70	91.60
HOG (reference)	23.92	95.55	94.60

B: Experiment 2: Effect of Training Set Size

Fig. 3 shows the relationship between training set size, J2 criterion value, and SVM accuracy using HOG features. As the training set grows from 5,000 to 60,000 samples, recognition accuracy increases steadily from 94.15 percent to 96.75 percent, a gain of 2.60 percentage points. Notably, the J2 criterion value remains approximately stable across all four training sizes, varying only between 23.59 and 24.42. This is an important finding: it confirms that the J2 criterion value is largely independent of training data volume for a fixed number of classes, and that the accuracy improvement from additional data arises primarily from better classifier boundary estimation rather than from improved compression quality.

Fig. 3: EFFECT OF TRAINING SET SIZE ON J2 VALUE AND SVM ACCURACY USING HOG FEATURES. Accuracy rises from 94.15 PERCENT AT 5,000 SAMPLES TO 96.75 PERCENT AT 60,000 SAMPLES, WHILE THE J2 VALUE REMAINS APPROXIMATELY STABLE, INDICATING THAT COMPRESSION QUALITY IS LARGELY DATA-VOLUME INDEPENDENT.

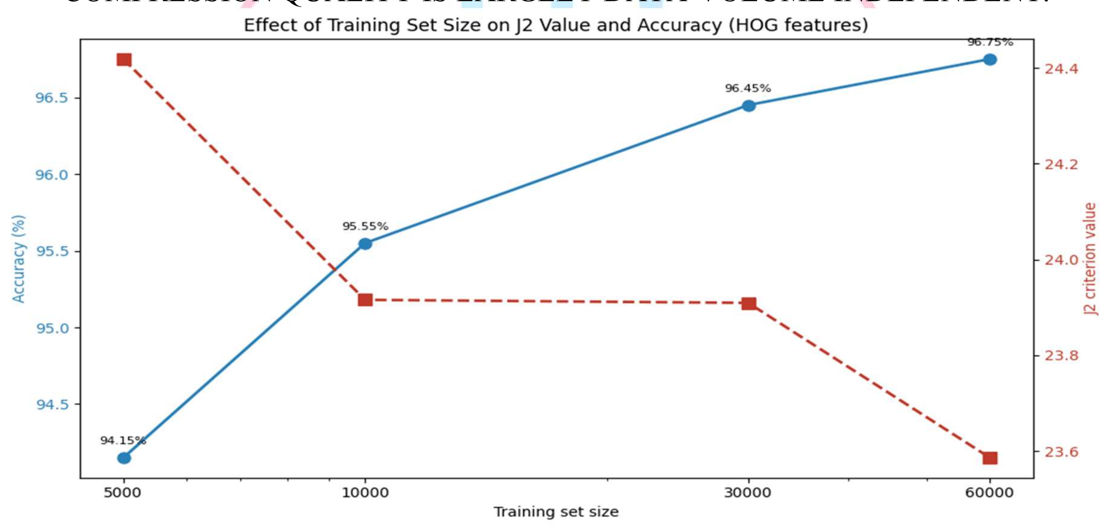

Table 4: Experiment 2 results: J2 value and SVM accuracy versus training set size (HOG features).

Training Size	J2 Value	SVM Accuracy (%)
5,000	24.42	94.15
10,000	23.92	95.55
30,000	23.91	96.45
60,000	23.59	96.75

C. Experiment 3: Multi-Descriptor Fusion

Fig. 4 compares the accuracy of the individual and fused descriptors under J2 compression and SVM classification. The SIFT Bag of Words descriptor alone achieves only 67.60 percent accuracy, reflecting the sparsity of keypoints detected on small 28 by 28 digit images. However, when combined with the other descriptors, SIFT contributes complementary information. The baseline LBP+HOG fusion achieves 95.95 percent, while adding SIFT to form the triple LBP+HOG+SIFT fusion yields a marginally lower 95.75 percent, indicating that baseline LBP and SIFT overlap in the information they contribute. The full fusion of SP-LBP+HOG+SIFT achieves the highest accuracy of 97.30 percent with the largest J2 value of 38.17, an improvement of 1.35 percentage points over the previous best result of 95.95 percent from the baseline study.

JOURNAL OF ENGINEERING AND COMPUTATIONAL INTELLIGENCE REVIEW (JECIR)
ISSN (Online): 3105-4811 ISSN (Print): 3106-969X
Volume 4, Issue 1, 2026
https://jecir.com/

Fig. 4: FEATURE FUSION COMPARISON UNDER J2 COMPRESSION AND SVM CLASSIFICATION. THE FULL FUSION OF SP-LBP, HOG, AND SIFT BAG OF WORDS ACHIEVES THE HIGHEST ACCURACY OF 97.30 PERCENT, DEMONSTRATING THAT THE THREE DESCRIPTORS CARRY COMPLEMENTARY DISCRIMINATIVE INFORMATION.

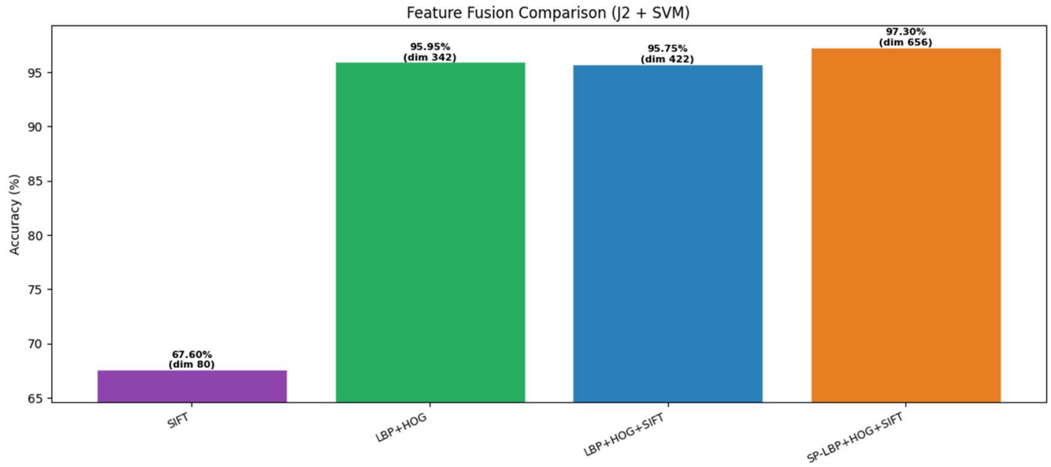

Table 5 summarises all fusion configurations. The monotonic increase of the J2 value from 26.20 for the baseline fusion to 38.17 for the full fusion confirms that the addition of spatially-encoded texture and keypoint information progressively increases class separability. The full fusion result of 97.30 percent validates the central hypothesis that the three descriptors are complementary, with SP-LBP capturing spatially-arranged texture, HOG capturing global stroke geometry, and SIFT capturing scale-invariant keypoint structure.

Table 5: Experiment 3 results: accuracy and J2 value for individual and fused feature descriptors

Configuration	Dim	J2	SVM (%)	LDA (%)
SIFT-BoW	80	6.45	67.60	64.45
LBP+HOG (baseline)	342	26.20	95.95	95.60
LBP+HOG+SIFT	422	30.73	95.75	95.85
SP-LBP+HOG+SIFT	656	38.17	97.30	97.00

D. Discussion

The results in the three experiments are consistent in several ways: First, the J2 criterion value is a reliable indicator of the recognition accuracy for all the feature configurations: higher J2 value is directly proportional to the recognition accuracy ranging from 56.25 percent for baseline LBP to 97.30 percent for the full fusion. Second, spatial information is important for texture descriptors, as is seen by the dramatic improvement from global LBP to spatial pyramid LBP. Third, feature fusion is helpful if the descriptors describe complementary information, but not if the descriptors have overlapping information, such as LBP+HOG and LBP+HOG+SIFT. In all experiments, SVM is marginally better than LDA, with the difference decreasing as feature space quality increases, as this was found in the baseline study to be nearly linearly separable in the J2-compressed subspace.

One of the drawbacks of this current study is the lack of evaluation on the EMNIST data, as was suggested in the baseline paper under the fourth future-work direction. During experimentation, the EMNIST-Digits partition has been found to be unavailable in the public repository and is thus left as an open task for future research. The other three directions have all been resolved, and the fusion configuration has been used to achieve a new best accuracy of 97.30 percent in the J2 compression framework.

IV. CONCLUSION AND FUTURE WORK

This paper extended our previous work on feature compression for handwritten digit recognition based on the J2 criterion, outlining three directions for future work. A Spatial Pyramid Local Binary Pattern (SPLBP) descriptor was proposed and found to improve texture-only recognition accuracy from 56.25 percent to 92.70 percent by preserving the spatial structure of local texture, which is lost in a global histogram, thereby reducing most of the gap between texture and gradient-based features. The J2 criterion value is roughly equal across all cases, and a data-scaling study from 5,000 to 60,000 training samples showed that accuracy improved from 94.15 % to 96.75 %, with no significant effect on compression quality as the training data volume was varied for a fixed number of classes. The most accurate fusion result is the triple fusion of Spatial Pyramid Local Binary Pattern, Histogram of Oriented Gradients and Scale Invariant Feature Transform Bag of Words features which performed with highest accuracy of 97.30 percent with the highest J2 value of 38.17 which is 1.35 percent better compared to the best accuracy obtained previously, thereby confirming the complementary nature of the three descriptors used. The from scratch J2 compression and the Support Vector Machine along with the Linear Discriminant Analysis classifier were kept the same throughout to allow for a controlled comparison.

There are a number of directions for future research. The most immediate is testing the proposed framework on the EMNIST dataset to verify generalisation to a different population of writers, which could not be performed in the present study due to the dataset's limited availability. Additional directions are the exploration of learned pooling strategies which can be used instead of the fixed spatial pyramid grid, incorporating denser keypoint sampling to boost the SIFT descriptor on small images further, extending the J2 framework to a kernelised formulation, which can model more non-linear class boundaries, and the application of the complete pipeline to more challenging real-world handwriting datasets with additional sources of difficulty, including font variation, background texture and stroke degradation.

V. REFERENCES

[1] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," Proc. IEEE, vol. 86, no. 11, pp. 2278–2324, 1998.

[2] R. Plamondon and S. N. Srihari, "Online and offline handwriting recognition: A comprehensive survey," IEEE Trans. Pattern Anal. Mach. Intell., vol. 22, no. 1, pp. 63–84, 2000.

[3] A. Baldominos, Y. Saez, and P. Isasi, "A survey of handwritten character recognition with MNIST and EMNIST," Appl. Sci., vol. 9, no. 15, p. 3169, 2019.

[4] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nat. Mach. Intell., vol. 1, no. 5, pp. 206–215, 2019.

[5] S. T. Hasan, "Natural language processing for employee relations case management: Analyzing workplace complaints, grievances, and investigation narratives in healthcare organizations," J. Comput. Syst. Eng. Insights, vol. 1, no. 1, pp. 11–24, 2023.

[6] U. Iqbal, "AI-powered supplier risk intelligence: Predicting financial and geopolitical supply chain disruptions in U.S. critical industries," J. Eng. Comput. Intell. Rev., vol. 3, no. 2, pp. 173–193, 2025.

[7] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," IEEE Trans. Pattern Anal. Mach. Intell., vol. 19, no. 7, pp. 711–720, 1997.

[8] R. Parveen, "An explainable AI framework for ethical fraud prevention in U.S. federal welfare programs," Int. J. Appl. Math., vol. 38, no. 11S, pp. 1425–1435, 2025.

[9] S. T. Hasan, "Predictive modeling for labor law risk detection in Medicaid funded home care programs in New York: Forecasting overtime exposure, payroll errors, and workforce compliance gaps," Int. J. Bus. Integr. Technol. Adv., vol. 1, no. 1, pp. 1–10, 2025.

[10] S. Akter, "Exploring the role of artificial intelligence in food waste reduction: Implications for public health, social behavior, and sustainable communities," J. Comput. Syst. Eng. Insights, vol. 1, no. 1, pp. 25–36, 2023.

[11] M. M. Islam, "AI-driven fraud detection and prevention using human behavior analysis to enhance U.S. social and financial security," Int. J. Appl. Math., vol. 38, no. 8S, pp. 861–871, 2025.

[12] U. Iqbal, "AI-driven predictive maintenance for U.S. smart manufacturing: Deep learning models for equipment failure prediction and operational resilience," J. Eng. Comput. Intell. Rev., vol. 3, no. 1, pp. 114–138, 2025.

[13] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray, "Visual categorisation with bags of keypoints," in Proc. ECCV Workshop Statistical Learning in Computer Vision, 2004, pp. 1–22.
[14] I. Kate, "Marketing innovation in the US e-commerce sector: Strategies for success in the evolving digital environment," Int. J. Bus. Integr. Technol. Adv., vol. 1, no. 2, pp. 1–7, 2025.
[15] A. Gupta et al., "A facile approach to recycling used facemasks for high-performance energy-storage devices," in Specialty Polymers. Boca Raton, FL, USA: CRC Press, 2023, pp. 427–443.
[16] S. M. H. Shah, F. Amin, and A. Khan, "Cyber-resilient mobile edge computing: A deep neural approach for secure and efficient task offloading," Asian Bull. Big Data Manag., vol. 5, no. 1, pp. 200–215, 2025.
[17] I. A. Badhan, M. N. Hasnain, and M. H. Rahman, "Enhancing operational efficiency: A comprehensive analysis of machine learning integration in industrial automation," J. Bus. Insight Innov., vol. 1, no. 2, pp. 61–77, 2022.
[18] U. Imtiaz, B. Ahmad, M. H. Sajid, Q. Abbas, M. A. Qureshi, S. Rasheed, and A. Khan, "An integrated machine learning framework for structural health monitoring of bridges: A case study on Soan Bridge," Asian Bull. Big Data Manag., vol. 5, no. 2, pp. 194–207, 2025.
[19] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.
[20] G. Bradski, "The OpenCV library," Dr. Dobb's J. Softw. Tools, vol. 25, no. 11, pp. 120–125, 2000.
[21] D. Impedovo and G. Pirlo, "Machine learning for handwriting recognition," Entropy, vol. 20, no. 11, p. 821, 2018.
[22] F. Amin, M. A. But, I. Amin, and A. Khan, "The tokenized business marketplace: A blockchain and AI-powered framework for democratizing business ownership and investment," Int. J. Bus. Manag. Sci., vol. 5, no. 4, pp. 318–328, 2024.
[23] U. Imtiaz, "Investigating the impact of phishing attacks on organizational cybersecurity posture," Spectrum Eng. Sci., pp. 1758–1780, 2025.
[24] I. A. Badhan, M. N. Hasnain, and M. H. Rahman, "Advancing operational efficiency: An in-depth study of machine learning applications in industrial automation," Policy Res. J., vol. 1, no. 2, pp. 21–41, 2023.
[25] S. Bhardwaj et al., "Bimetallic Co–Fe sulfide and phosphide as efficient electrode materials for overall water splitting and supercapacitor," Discover Nano, vol. 18, no. 1, p. 59, 2023.
[26] X. Chen, L. Jin, Y. Zhu, C. Luo, and T. Wang, "Text recognition in the wild: A survey," ACM Comput. Surv., vol. 54, no. 2, pp. 1–35, 2021.
[27] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," IEEE Intell. Syst., vol. 13, no. 4, pp. 18–28, 1998.