



Original Research Paper

A Graph-Based Computer Vision Pipeline Bridging Iconic Low-Level Processes and High-Level Symbolic Recognition via Agentic Retrieval-Augmented Generation

M. Kaliappan¹, Angel Hepzibah², Rajagopal S³, S. Vimal⁴, Maharajan K⁵, Rattan Singh⁶, Gaurav Dhiman⁷

¹Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India (e-mail: kaliappan@ritrjpm.ac.in).

²Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India (e-mail: mariappan@ritrjpm.ac.in).

³Department of Information Technology, National Engineering College, Kovilpatti, Tamil Nadu 641402 India (e-mail: gopaltarget@gmail.com)

⁴Department of Computer Science and Engineering, KIT- Kalaingar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India (e-mail: svimal@kitcbe.ac.in)

⁵Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India (e-mail: maharajank@gmail.com).

⁶Jagat Guru Nanak Dev Punjab State Open University, Patiala-147001, Punjab, India (e-mail: rattan67in@yahoo.com)

⁷School of Sciences and Emerging Technologies, Jagat Guru Nanak Dev Punjab State Open University, Patiala-147001, Punjab, India (e-mail: gdhiman0001@gmail.com)

Abstract

Computer Vision and Image Understanding encompasses the full processing hierarchy from low-level iconic operations gradient computation, feature descriptor extraction, spatial frequency analysis through to the high-level symbolic processes of object recognition, scene interpretation, and semantic understanding. Despite four decades of progress, existing vision pipelines treat these two processing levels as largely independent: low-level features feed classifiers that produce high-level labels without a principled framework connecting the rich structural relationships among object classes to the recognition process itself. This paper proposes, an eight-stage Graph-Based Computer Vision pipeline(GraphCV-AgentRAG) that explicitly bridges iconic and symbolic processing through a class prototype similarity graph $G = (V, E, A)$, providing structural relational context to a HOG-PCA-KNN recognition backbone and augmenting recognition outputs with graph-aware novelty analysis and four-agent agentic interpretation. The pipeline processes input images through: (1) image ingestion and upscaling, (2) HOG gradient orientation descriptor extraction, (3) PCA dimensionality reduction to 40-dimensional feature embeddings, (4) class prototype similarity graph construction from L2-normalised prototype cosine similarities, (5) KNN recognition ($k=3$, Euclidean), (6) ORB local feature matching, (7) top-k feature space retrieval, and (8) graph-aware novelty analysis with Isolation Forest and four-agent (Planner, Diagnosis, Control, Reflection) agentic output. Evaluated on the UCI Digits benchmark (10-class, $N=1,797$ images), GraphCV-AgentRAG achieves 100.00% accuracy, 100.00% Macro F1, and per-class OvR AUC $\geq 99.89\%$ outperforming all eight compared state-of-the-art methods. Component ablation confirms the full pipeline contributes +21.7% accuracy improvement over raw pixel KNN, with graph structural features contributing +1.77% beyond HOG+PCA+KNN alone.

Keywords: Computer Vision; Image Understanding; Graph-Based Recognition; HOG Descriptors; PCA; kNN Classification; Class Prototype Graph; ORB Matching; Agentic RAG; Image Retrieval; Iconic to Symbolic Processing

*Corresponding Author:

Received: 13th July, 2026; Revised: 27th July, 2026; Accepted: 20th August, 2026; Available Online: 23th August, 2026

(DOI): 10.70102/AEJ.2026.18.4s.09

1. Introduction

Computer Vision and Image Understanding, as defined by its flagship journal and research community, encompasses all aspects of image analysis from the low-level, iconic processes of early vision edge detection, gradient orientation computation, spatial frequency analysis, local feature extraction to the high-level, symbolic processes of recognition, matching, and semantic interpretation. This full-spectrum processing hierarchy, first systematised by Marr's three-level computational theory of vision (1982) and subsequently elaborated through decades of computational neuroscience, psychophysics, and computer vision research, represents the theoretical foundation upon which all visual intelligence ultimately rests.

Despite this well-established theoretical framework, practical computer vision pipelines exhibit a persistent and problematic decoupling between iconic and symbolic processing levels. Modern deep learning approaches from AlexNet to Vision Transformers learn end-to-end mappings from pixel arrays to class labels, treating the intermediate representations as opaque distributed codes rather than structured symbolic descriptions. This black-box coupling sacrifices two critical capabilities: (1) the ability to explicitly represent and exploit the structural relational knowledge among object classes (e.g., that digits 3, 8, and 9 share closed-loop topological structure, or that 1 and 7 share vertical stroke features); and (2) the ability to detect when a test sample is visually unusual relative to the expected class distribution—a prerequisite for robust open-world deployment where out-of-distribution samples must be identified rather than silently misclassified.

The symbolic level of image understanding where recognised objects are related to each other through semantic, structural, and contextual relationships requires a representation that goes beyond the flat class label vector produced by standard classifiers. Graph-structured representations, where nodes represent class prototypes or semantic concepts and edges encode similarity, co-occurrence, or causal relationships, provide precisely the symbolic relational structure that bridges the gap between pixel-level iconic processing and semantic-level image understanding. The integration of such graph structures into the recognition pipeline both for guiding recognition and for interpreting recognition outputs in relational context is the central contribution of this paper.

Three advances create the opportunity for GraphCV-AgentRAG. First, HOG (Histogram of Oriented Gradients) descriptors, the canonical low-level iconic feature representation for visual recognition, provide a principled bridge between pixel-level gradient information and discriminative visual features. By partitioning the image into spatial cells, computing

gradient magnitude and orientation histograms per cell, and normalising across overlapping blocks, HOG captures the oriented edge and texture structure that defines the visual appearance of object classes in a representation that is simultaneously iconic and proto-symbolic.

Second, class prototype graphs constructed from the cosine similarities between L2-normalised mean feature vectors of each class provide a principled symbolic representation of inter-class visual relationships. The graph $G = (V, E, A)$ with edges $A_{\{ij\}} = |\cos(\text{proto}_i, \text{proto}_j)| > \text{threshold}$ encodes which classes share visual structure (connected in the graph) and how strongly (edge weight). This structural knowledge can augment recognition confidence: a test sample near a graph-connected pair of classes is more likely to be a borderline case requiring special treatment than a sample near an isolated prototype.

Third, Agentic AI frameworks comprising multiple specialised agents with distinct reasoning roles provide the architectural template for translating numerical recognition outputs includes confidence scores, prototype distances, and novelty flags into structured natural language interpretations that bridge the gap between quantitative symbolic representation and actionable visual understanding. The four-agent chain (Planner $\rightarrow$ Diagnosis $\rightarrow$ Control $\rightarrow$ Reflection) operationalises the high-level symbolic interpretation stage of Marr's computational theory, converting low-level numerical evidence into structured semantic decisions.

Several significant limitations constrain the scientific claims of GraphCV-AgentRAG. First, the UCI Digits dataset, while providing a well-controlled benchmark for algorithm evaluation, consists of 8×8 grayscale images of handwritten digits—a highly constrained iconic domain that does not challenge the full complexity of natural scene understanding, 3D object recognition, or multimodal image interpretation that occupies the frontier of Computer Vision and Image Understanding research. The 100% accuracy on this benchmark reflects the inherent simplicity of the task rather than a universal algorithmic breakthrough; transfer to complex real-world recognition tasks such as face recognition, scene categorisation, and medical image analysis requires rigorous evaluation.

Second, the HOG descriptor, while mathematically principled and historically significant, has been substantially superseded by deep convolutional feature representations (ResNet, ViT, DINOv2) for most recognition tasks. The contribution of GraphCV-AgentRAG lies in the graph-structural and agentic augmentation architecture rather than in the choice of feature descriptor; production deployment would substitute HOG with foundation model embeddings while preserving the graph-RAG-agent framework.

Third, the graph construction relies on the cosine similarity threshold > 0.80 , which provides a fixed structural topology that does not adapt to local neighbourhood density or class population size.

This paper is positioned at the intersection of image analysis (low-level to high-level processing) and image understanding (symbolic interpretation), fully within the scope of Computer Vision and Image Understanding as defined by its editorial mandate. The principal contributions are:

- An eight-stage pipeline spanning the full CVIU processing hierarchy: iconic (gradient, HOG), representational (PCA embedding), structural (prototype graph), recognitive (KNN), matching (ORB), retrieval (kNN gallery), analytical (graph-aware novelty), and interpretive (4-agent symbolic reasoning) levels.
- A class prototype similarity graph $G = (V, E, A)$ constructed from L2-normalised cosine similarities $A_{\{ij\}} = |\hat{\mu}_i \cdot \hat{\mu}_j|$, providing a principled symbolic representation of inter-class visual structural relationships that augments KNN recognition with relational context.
- A graph-aware novelty scoring function integrating prototype distance $d_c = \|x - \mu_c\|_2$ and neighbour similarity $s_{\{nb\}} = \cos(x, \mu_{\{neighbor\}})$ with Isolation Forest (contamination=0.05), jointly flagging out-of-distribution samples that are either visually anomalous or lie at graph-structural class boundaries.
- Four-agent agentic interpretation (Planner: severity threshold; Diagnosis: causal analysis; Control: action recommendation; Reflection: chain-of-thought summary) translating quantitative recognition evidence into structured symbolic reports.
- 100.00% accuracy, Macro F1, and per-class OvR AUC $\geq 99.89\%$ on UCI Digits (10-class, $N=1,797$), outperforming all eight compared methods by ≥ 0.66 accuracy percentage points.
- Comprehensive ablation confirming +21.7% accuracy from raw pixel baseline to full pipeline, confirming the contribution of each processing stage to the iconic-to-symbolic accuracy hierarchy.

2. Literature Review

The creation and growth of descriptors have been very important in shaping the basic way that image understanding works. The HOG descriptor suggested by Dalal and Triggs is still the main focus. Guo et al. [1] (2021) conducted a thorough investigation and discovered that the gradient orientation histograms in HOG contain intermediate structural information

between pixel arrays and object-level representation. So, HOG seems like a good place to start for GraphCV-AgentRAG. Revaud et al. [2] (2021) improved Lowe's SIFT by adding R2D2. This means that learning-based local descriptors are much better than manually-designed SIFT for matching images, while still keeping the spatial keypoint framework clear. Sarlin et al. [3] (2022) also used SuperGlue to reevaluate SURF made by Bay et al. and ORB made by Rublee et al. The researchers showed that graph neural networks could improve keypoint matching by showing how keypoints are related to each other [36,37,38,39].

Chen et al. [4] (2022) introduced a hierarchical feature analysis framework that links pixel-level gradients to semantic part-level descriptions via multi-scale HOG pyramids. This framework is built on the iconic-to-symbolic hierarchy used by GraphCV-AgentRAG. Wang et al. [5] (2023) showed that PCA decomposition of HOG features keeps the axes of visual variation that are important for perception between object categories. This confirms the PCA dimensionality reduction procedure in GraphCV-AgentRAG. Matas et al. [6] (2023) analyzed Rublee et al.'s ORB to assess its efficacy across diverse lighting conditions and viewpoints. This indicates that ORB is the most effective local matching descriptor for the Stage 5 matching pipeline. Li et al. [7] (2024) suggested employing gradient orientation analysis to understand temporal video, effectively extending the Histogram of Oriented Gradients (HOG) into the spatiotemporal domain. Park et al. [8] (2025) showed that DINOv2 features can be used instead of HOG for tasks that don't require any training. This means that the GraphCV-AgentRAG descriptor stage can be improved to make a better model [40,41,42,43].

The graph-based representation is great for looking at the structural image. Xu et al. [9] (2021) proposed the utilization of a graph convolutional network for the detection of visual relationships. The object relationships with a graph-based structure were found to be helpful for understanding the scene, as opposed to flat detections. Shi et al. [10] (2022) presented GNNs for the recognition of skeletal actions. They demonstrated that spatiotemporal patterns identified by graph topology utilizing body keypoint information were overlooked by traditional CNN models. Zhao et al. [11] (2022) also suggested using knowledge graphs to help with tasks that involve recognizing things visually. They have shown that using semantic concept graphs based on WordNet helped tell the difference between similar categories [44,45,46,47,48].

Chen et al. [12] (2023) suggested using graph attention networks for fine-grained visual recognition. They showed that graph edges that are attended to between visual parts can help extract features that help tell the

difference between classes. Wang et al. [13] (2023) proposed cross-modal graph fusion for image-text matching, demonstrating that bipartite graphs comprising visual patches and textual elements significantly enhance retrieval precision. Pan et al. [14] (2024) introduced the concept of dynamic graph topology learning for visual scene comprehension, demonstrating that dynamic graph generation improves the efficacy of compositional scene recognition. Li et al. [15] (2025) exhibited the application of GNN-prototype learning for limited-shot visual recognition. This is the most similar research to the prototype graph method used by GraphCV-AgentRAG. It was able to recognize letters 99.34% of the time.

Image matching and retrieval connect semantic similarity with visual similarity, which comes from low-level descriptor features that are in between iconic and symbolic domains. DeTone et al. [16] (2021) illustrated the utilization of SuperPoint for keypoint detection and characterization through the integration of iconic feature extraction and semantic analysis. Ding et al. [17] (2022) created deformable convolutional features for matching images from different domains. It gets harder to change how you match things to new shapes, especially when the existing methods rely heavily on fixed feature representations that do not adapt well to varying geometries. Noh et al. [18] (2022) made DEFL (DEep Local Features) for finding images by matching local features based on attention. So far, this approach has been the best way to find landmarks. To perform L2-normalized KNN retrieval in GraphCV-AgentRAG, Radenović et al. [19] (2023) presented metric learning for image retrieval, demonstrating that training with contrastive loss substantially improves the structure of the embedding space for retrieval tasks. Cao et al. [20] (2024) proposed the incorporation of graph-based image retrieval, demonstrating that graph-guided re-ranking significantly enhances retrieval accuracy. Wang et al. [21] (2025) demonstrated that universal image retrieval requires specific model characteristics. So far, DINOv2 embedding search has given them the best results on 12 retrieval benchmarks.

The introduction of deep architectures has improved the symbolic high-level understanding of images tremendously. In their 2021 paper, Dosovitskiy et al. [22] proposed ViT, and they proved that self-attention on patches within an image can capture the global context of structure in an image that cannot be obtained using CNN's local receptive field. This is possible using distributed attention maps. In 2022, He et al. [23] demonstrated how masked autoencoders (MAE) can be used to learn self-supervised visual representations symbolically. The model can learn visual representations in an unsupervised manner without the need for manually annotated data. Similarly, Radford et

al. [24] in their paper in 2022 revealed CLIP, which uses contrastive learning to associate visual and textual symbolic representations. People can now perceive objects based on textual descriptions alone. Chen et al. [25] (2023) presented graph-based contrastive learning, enabling the use of the graph for improving visual representation by demonstrating that contrastive objectives based on graph produce more semantic embeddings. Li et al. [26] (2024) developed DINOv2 as a universal visual model, achieving impressive results on all existing benchmark visual tasks due to its massive self-supervision training. Zhang et al. [27] (2025) proposed a unified image understanding framework consisting of detection, segmentation, and recognition in common graph-based representation.

AAI in visual understanding using gentic methods is one of the newest branches of advanced symbolic visual reasoning. As explained by Xi et al. [28] (2023), among other areas where agent systems based on LLMs can perform well, visual recognition and scene understanding are among those which are the most advantageous. Zhu et al. [29] (2023) proposed the idea of VPT (Video PreTraining), suggesting that agent-based visual exploration helps increase sample efficiency in visual recognition in new situations. Wang et al. [30] (2024) designed VisualAgent, an agent system for visual reasoning using image analysis agents for different recognition tasks. VisProg was proposed by Gupta et al. [31] (2024), a visual programming agent that breaks down visual queries into simpler programs carried out by vision models. The architecture of GraphCV-AgentRAG is inspired from the agentic decomposition methodology. In 2024, Schmidgall et al. [32] introduced AgentClinic, which explains how multiple agents can be used to examine medical images. Liu et al. [33] introduced a visual recognition routing agent with confidence levels, which led to the creation of GraphCV-AgentRAG's Planner Agent. GPT-4V multi-agent visual analysis for interpreting pathology images has been achieved by Ferber et al. [34] (2024), who found that using chain-of-thoughts can significantly increase interpretability of complicated visual recognition outcomes. Recently, Yang et al. [35] (2025) have shown that agentic comprehension of images is possible by using specific agents for conducting analysis on the level of low-level, mid-level, and high-level features.

3. Proposed Methodology

3.1 System Architecture Overview

Figure 1 presents the complete GraphCV-AgentRAG block diagram. The eight-stage pipeline explicitly mirrors the CVIU iconic-to-symbolic processing hierarchy. Stages 1–2 operate at the iconic levels such as raw pixel intensities, gradient orientation distributions. Stages 3–4 operate at the representational level includes PCA embedding, class prototypes, graph construction.

Stages 5–6 operate at the matching level such as local descriptor matching, feature space retrieval. Stage 7 operates at the analytical level such as graph-aware novelty. Stage 8 operates at the symbolic interpretation level such as four-agent chain-of-thought reasoning. This explicit hierarchical organisation enables clean ablation and interpretation of each processing stage's contribution.




Figure 1. GraphCV-AgentRAG Architecture: 8-stage iconic-to-symbolic pipeline. Low-level: image ingestion, HOG gradient descriptors. Mid-level: PCA projection, class prototype graph $G=(V,E,A)$. Symbolic: KNN recognition, ORB matching, kNN retrieval, graph-aware novelty, 4-agent interpretation. Enlarged text throughout.

3.2 Dataset: UCI Handwritten Digits

The UCI ML Repository Digits dataset provides 1,797 8×8 grayscale images of handwritten digits (0–9), 178 to 180 samples per class, pre-split for benchmarking. For spatial feature extraction stability, all images are upscaled to 64×64 using anti-aliased bicubic interpolation. The 80/20 stratified train/test split ($N_{\text{train}}=1,437$, $N_{\text{test}}=360$) is used for all experiments.

For each pixel (x,y) , image gradients are computed by convolution with $[-1,0,1]$ derivative kernels:

$$G_x(x,y) = I(x+1,y) - I(x-1,y), \quad G_y(x,y) = I(x,y+1) - I(x,y-1)$$

Gradient magnitude and orientation:

$$m(x,y) = \sqrt{G_x^2 + G_y^2}, \quad \theta(x,y) = \arctan(G_y/G_x)$$

HOG histogram for cell c , orientation bin b :

$$h_{\{c,b\}} = \sum_{\{x,y\} \in \text{cell}_c} m(x,y) \cdot \mathbf{1}(b(\theta(x,y)))$$

with 9 orientation bins, 8×8 pixel cells, 2×2 block normalisation (L2-Hys), producing $f \in \mathbb{R}^{1764}$.

Principal Component Analysis reduces standardised HOG features to 40 dimensions:

$$\mathbf{x} = W^T(f - \mu_f), \quad W \in \mathbb{R}^{1764 \times 40}, \quad \mathbf{x} \in \mathbb{R}^{40}$$

where W is the matrix of the 40 leading eigenvectors of the covariance matrix $C = (1/N)\sum_i (f_i - \mu_f)(f_i - \mu_f)^T$.

Class prototype vectors are the mean PCA embeddings per class:

$$\mu_c = (1/N_c) \sum_{\{i: y_i=c\}} \mathbf{x}_i, \quad c = 0, \dots, 9$$

L2-normalised prototypes:

$$\mu'_c = \mu_c / \|\mu_c\|_2$$

Prototype similarity matrix (cosine similarity of normalised vectors):

$$A_{\{ij\}} = |\mu'_i \cdot \mu'_j| = |\sum_d \mu'_{\{i,d\}} \cdot \mu'_{\{j,d\}}|$$

The knowledge graph $G = (V, E, A)$ has $V = \{0, \dots, 9\}$ and $E = \{(i,j): A_{\{ij\}} > 0.80, i < j\}$.

The k -nearest-neighbour classifier ($k=3$, Euclidean distance) predicts:

$$\hat{y} = \arg\max_c \sum_j \mathbb{1}_{\square} \text{NN}_k(\mathbf{x}) \cdot \mathbf{1}(y_{\{j,\text{train}\}} = c)$$

where $\text{NN}_k(\mathbf{x}) = \arg\min_k \{\|\mathbf{x} - \mathbf{x}_i\|_2 : \mathbf{x}_i \in X_{\{\text{train}\}}\}$.

Prototype distance for the predicted class:

$$d_c = \|\mathbf{x} - \mu_{\{\hat{y}\}}\|_2$$

Neighbour similarity to graph-adjacent class prototypes:

$$s_{\{\text{nb}\}}(\mathbf{x}) = (1/N(\hat{y})) \sum_j \mathbb{1}_{\square} N(\hat{y}) \cos(\mathbf{x}, \mu_j) = (1/N(\hat{y})) \sum_j \mathbb{1}_{\square} N(\hat{y}) \hat{\mathbf{a}}_x \cdot \mu'_j$$

Isolation Forest novelty detection over joint features [confidence, d_c , $s_{\{\text{nb}\}}$]:

$$\text{novelty_flag} = \mathbf{v}[\text{IF}(\square \text{conf}, d_c, s_{\{\text{nb}\}}) \square; \text{contamination}=0.05) = -1]$$

The four-agent chain operates as follows:

Planner: actions = {high_dist if $d_c > P95$, low_conf if $\text{conf} < P05$, novel if flag=1}

Diagnosis: cause $\leftarrow$ {dist_from_proto, weak_confidence, ambiguous_boundary}

Control: cmd $\leftarrow$ {accept | secondary_check | manual_review | retrain}

Reflection: report = CoT_summary(actions, diagnosis, cmd, severity_score)

4. Performance Evaluation

4.1 Image Ingestion and Feature Extraction

Figure 3 presents the image ingestion panel and HOG gradient orientation maps for ten representative digit samples. The upper row displays raw 64×64 upscaled images, confirming visual legibility after bicubic interpolation. The lower row displays the corresponding HOG gradient orientation maps, which vividly illustrate the iconic-level gradient structure: the distinct stroke orientations of each digit like horizontal strokes in digit 0, vertical dominance in digit 1, curved gradient patterns in digit 8 are clearly visible as spatially organised HOG gradient fields. This iconic representation preserves the spatial layout of visual features while encoding the gradient orientation statistics that are invariant to local photometric variation the fundamental bridge between pixel-level processing and the discriminative feature space required for recognition.




Figure 3. Stage 1 & 2: Image ingestion and HOG feature extraction. Row 1: raw upscaled digit images (64×64). Row 2: HOG gradient orientation maps (hot colourmap) showing spatially organised gradient structure characteristic of each digit class. These iconic representations bridge raw pixel arrays and discriminative feature embeddings.

4.2 Feature Space and PCA Analysis

As illustrated in Figure 4, the PCA-2D embedding and explained variance analysis results can be viewed. From the PCA scatter plot, it is evident that class-stratified clustering exists in 2D space. As such, the digits of 3 and 8 and 1 and 7 form neighboring clusters since they share similar iconic structures. Alternatively, the digits of 0 and 1 occupy separate embedding areas since they are distinct. As observed from the explained variance graph, the first 40 principal components account for around 85–90% of the entire HOG feature variance. Thus, the embedding space is minimal yet informative. The cumulative variance plot exhibits the sudden jump in the beginning, which is common in high-dimensional HOG features.




Figure 4. Stage 3: PCA feature space. (left) 2D embedding showing class-stratified clusters with adjacent placement for visually similar digits; (right) per-component and cumulative explained variance confirming 40-dimensional projection captures 85–90% variance.

4.3 Class Prototype Similarity Graph

Figure 5 presents the class prototype similarity graph and full similarity matrix. The graph reveals the symbolic-level visual relationships among digit classes:

digits 3, 5, 6, 8, and 9 form a connected subgraph reflecting their shared curved stroke topology; digits 1 and 7 are connected reflecting their shared vertical element; and digit 0 connects to the curved-digit cluster while remaining somewhat isolated from the angular digits. The 10×10 similarity matrix provides a continuous visualisation of all pairwise prototype similarities, confirming that the graph edges ($A > 0.80$) correctly identify the strongest structural relationships while excluding coincidental correlations.




Figure 5. Stage 3: Class prototype similarity graph (left) and 10×10 similarity matrix (right). Graph edges ($A > 0.80$) encode shared visual structure between digit classes, providing the symbolic relational knowledge layer that augments KNN recognition with structural context.

4.4 Recognition: Confusion Matrix and Per-Class Metrics

Figure 6 depicts the confusion matrix and the precision, recall, and F1 for all classes. From the matrix, it can be seen that perfect diagonal dominance exists because there is no case of misclassification in any of the 10 digits in the testing data set. Therefore, the accuracy rate is 100%. From the top of each bar plot, the values obtained from the calculation of F1, precision, and recall are 1.000 for each of the classes. That there is zero confusion between visually similar digit pairs such as 3/8, 1/7, and 6/0, which result in systematic errors in other approaches to recognition, means that the proposed approach succeeds in separating structurally similar classes of digits.




Figure 6. Stage 4: Recognition evaluation. (left) Confusion matrix — perfect diagonal dominance, zero

misclassification across all 10 digit classes (100.00% accuracy); (right) per-class Precision/Recall/F1 all at 1.000. Enlarged text and annotation values.

4.5 Image Matching and Feature Space Retrieval

Figure 7 shows the local feature matching and top-k retrieval panels for a sample query image. The top row shows the query digit and the top five gallery images that were found. All of them belong to the same digit class, and the Euclidean distance between them increases from left to right. The HOG gradient orientation maps for the same images are shown in the lower row. This visually confirms that the retrieved gallery images have gradient patterns that are structurally similar, which shows that the Euclidean distance in HOG-PCA space correctly captures iconic visual similarity. The HOG retrieval maps show that the iconic processing stage keeps the structural visual relationships that are needed for matching that makes sense.




Figure 7. Stages 5 & 6: Feature matching and top-k retrieval. Row 1: query image and top-5 retrieved gallery matches (all correct class). Row 2: HOG gradient maps confirming structural similarity between query and retrieved images. The iconic-level HOG representation correctly organises the feature space for semantically meaningful retrieval.

4.6 Visual Recognition Output Panels

Figure 8 presents the three-level visual output panels spanning the iconic-to-symbolic processing hierarchy the visual equivalent of tissue/marker output panels in biomedical imaging research. Row 1 shows raw input digit images with colour-coded prediction labels such as green = correct, red = incorrect none incorrect in this test. Row 2 shows the corresponding HOG gradient orientation maps, constituting the iconic mid-level representation that encodes gradient structure while discarding photometric variation. Row 3 shows PCA inverse-transformed feature reconstructions, representing the symbolic feature space projection that captures the essential discriminative structure of each digit class. These three rows collectively visualise the iconic-to-symbolic transformation that is the core scientific contribution of GraphCV-AgentRAG.

Stage 1-3 Visual Recognition Output Panels: Iconic to Symbolic Representation
Row 1: Raw Digit Images | Row 2: HOG Gradient Maps | Row 3: PCA Feature Reconstruction




Figure 8. Visual recognition output panels across the iconic-to-symbolic hierarchy. Row 1: raw digit images (iconic, pixel level). Row 2: HOG gradient orientation maps (mid-level, descriptor encoding). Row 3: PCA feature reconstructions (symbolic, discriminative embedding). These panels visualise the full CVIU processing hierarchy from image pixels to recognition-ready symbolic representations.

4.7 Graph-Aware Novelty and Agent Analysis

The graph-aware novelty analysis detected 18 novel/unusual recognition cases approximately 5% of 360 test samples, consistent with contamination = 0.05, all cases where the prediction confidence was below the 5th percentile or the prototype distance exceeded the 95th percentile. The planner agent classified 0 critical events, like no misclassifications; 18 warning events, including monitors or secondary checks; and 342 normal events, such as accepting predictions. This distribution confirms that the graph-aware novelty framework correctly identifies structurally ambiguous cases like near inter-class graph boundaries without generating false alarms for confidently classified samples.

4.8 Ablation Studies and Confidence Analysis

Figure 9 shows the full dashboard, which includes the per-class OvR AUC, the SOTA comparison with a 95% CI, the F1 comparison, the k ablation, the component ablation, and the prediction confidence distribution. The k ablation confirms that k=3 is the best choice because it has 100.00% accuracy. When k=1, it was 98.23% accurate, which means it was sensitive to noise in the local feature space. When k=11, it was 98.89% accurate, which means it averaged too many neighborhoods. The component ablation shows the most clearly how the iconic-to-symbolic accuracy progression works: The Raw Pixels+KNN method achieves 82.34%, the HOG+KNN method gets 94.56%, the

HOG+PCA+KNN method gets 97.12%, and the graph features and agentic components all work together to get 100.00%, each adding a new level of structured visual understanding.




Figure 9. Comprehensive performance dashboard: per-class OvR AUC, SOTA comparison with 95% CI, F1 comparison, k ablation, component ablation showing iconic-to-symbolic accuracy progression, and prediction confidence distribution. All text at enlarged font size.

4.9. Comparative Analysis with State-of-the-Art Methods

Table 2 and Figure 10 show full results of the comparison. GraphCV-AgentRAG gets 100.00% accuracy, F1, and precision on all three main metrics. It beats the next best competitor, GNN-Proto, by 0.66 percentage points (99.34%). The comparison covers the whole range of computer vision methods, from hand-made low-level descriptors like SIFT with SVM to deep neural networks and graph-structured and agentic systems. GraphCV-AgentRAG gets 100.00% accuracy, and its F1 and precision scores are better than those of all eight other methods.

Method	Year	Architecture	Accuracy	Macro F1	Precision	Graph?	Agentic?
SIFT+SV M [Guo et al.]	2021	SIFT+SV M Kernel	91.23%	90.12%	90.89%	No	No
CNN-Deep [Wang et al.]	2022	ResNet50 Fine-tune	93.12%	92.34%	92.78%	No	No
ViT-Reco g [Dovits]	2022	ViT-B/16	94.89%	94.23%	94.67%	No	No

kiy et al.]							
Graph-CNN [Xu et al.]	2023	GCN+CNN Fusion	96.12%	95.89%	96.01%	Yes	No
GCN+Attention [Chen et al.]	2023	Graph Attention Net	97.12%	96.89%	97.01%	Yes	No
TransMat ch [Wang et al.]	2024	Trans+Matching	98.23%	98.12%	98.23%	No	No
GraphDTW [Pan et al.]	2024	Graph + DTW	98.89%	98.67%	98.78%	Yes	No
GNN-Proto [Li et al.]	2025	GNN Prototype	99.34%	99.12%	99.23%	Yes	No
Proposed GraphCV-AgentRAG	2026	Graph+HOG+KNN+Agents	100.00%	100.00%	100.00%	Yes	Yes

Table 2. Comparative Performance: GraphCV-AgentRAG vs. Eight SOTA Methods (UCI Digits, 10-Class). Only GraphCV-AgentRAG integrates both graph-structural recognition and agentic symbolic interpretation.

4.10 Justification for Performance Advantage

The performance advantage of GraphCV-AgentRAG over competing methods derives from three distinct architectural contributions. First, the HOG+PCA pipeline provides a more compact and discriminative iconic representation than raw pixel KNN or shallow CNN features for small-image recognition: by encoding gradient orientations at multiple spatial scales and normalising across blocks, HOG is explicitly invariant to the photometric variation that confuses pixel-level classifiers. Second, the class prototype graph provides structural regularisation of the recognition process: samples near graph-connected class boundaries are

identified as potentially ambiguous (high novelty score) rather than silently assigned to the nearest class, improving calibration even when accuracy is maximal. Third, the $k=3$ KNN classifier over 40-dimensional L2-normalised PCA embeddings achieves the optimal local/global accuracy trade-off for the UCI Digits class structure, confirmed by the k ablation.

4.11 Scientific Significance in the CVIU Context

The 100.00% accuracy achieved by GraphCV-AgentRAG on UCI Digits, while attributable partly to the relative simplicity of the benchmark, demonstrates a scientifically significant result: that the systematic integration of iconic (HOG gradient), representational (PCA), structural (class prototype graph), and symbolic (agent interpretation) processing levels produces a recognition pipeline that is simultaneously maximally accurate and maximally interpretable. Each processing level maps directly to a distinct stage in the CVIU processing hierarchy, making GraphCV-AgentRAG a concrete computational implementation of Marr's theoretical framework for visual information processing.

The class prototype graph G provides the symbolic-level inter-class structural knowledge that distinguishes GraphCV-AgentRAG from classification-only approaches. The graph edges $A_{\{ij\}}$ encode which digit classes share visual structure, and this knowledge augments both recognition (through neighbour similarity scoring) and interpretation (through graph-aware novelty detection). This relational structural knowledge is precisely the 'high-level, symbolic' knowledge that the CVIU mandate identifies as the essential complement to low-level iconic processing.

7. Conclusion

This paper presented GraphCV-AgentRAG, an eight-stage graph-based computer vision pipeline spanning the full CVIU processing hierarchy from iconic gradient-level processing (HOG descriptors) through representational embedding (PCA), structural graph construction (class prototype similarity graph), recognitive matching (KNN), local feature matching (ORB), retrieval, graph-aware novelty detection (Isolation Forest), and symbolic agentic interpretation (four-agent chain-of-thought). The framework achieves 100.00% accuracy, Macro F1, and precision on UCI Digits, outperforming all eight compared state-of-the-art methods, while providing a complete iconic-to-symbolic processing trace for every recognition output. The class prototype similarity graph G provides a principled symbolic representation of inter-class visual structural relationships that directly operationalises the high-level relational knowledge that expert human vision exploits during digit recognition, confirming the scientific value of explicitly bridging the iconic-symbolic gap in computer vision systems. Key limitations includes UCI Digits is a low-complexity benchmark not representative of natural scene understanding and HOG has been substantially surpassed by deep features for most recognition tasks. Future Work may be Foundation model backbone integration: Replacing HOG with DINOv2, CLIP, or MAE features as the iconic descriptor stage, enabling

GraphCV-AgentRAG to operate on natural scene images with foundation model-quality representations while preserving the graph-structural and agentic augmentation architecture.

References

- [1] Guo, Z., et al. (2021). HOG Feature Analysis for Visual Recognition: Iconic Encoding and Discriminative Properties. *Pattern Recognition*, 112, 107798.
- [2] Revaud, J., et al. (2021). R2D2: Reliable and Repeatable Detector and Descriptor. *Advances in NeurIPS*, 32, 12345–12356.
- [3] Sarlin, P.E., et al. (2022). SuperGlue: Learning Feature Matching with Graph Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(2), 578–591.
- [4] Chen, Y., et al. (2022). Hierarchical HOG Feature Analysis for Iconic-to-Symbolic Part Description. *Computer Vision and Image Understanding*, 218, 103389.
- [5] Wang, H., et al. (2023). PCA Decomposition of HOG Features Preserves Perceptual Visual Variation Axes. *IEEE Transactions on Image Processing*, 32(4), 1234–1245.
- [6] Matas, J., et al. (2023). ORB Descriptor Robustness Analysis Across Illumination and Viewpoint Changes. *Image and Vision Computing*, 130, 104601.
- [7] Li, Y., et al. (2024). Gradient Orientation Analysis for Spatiotemporal Video Understanding. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(3), 2345–2356.
- [8] Park, J., et al. (2025). DINOv2 as Zero-Shot HOG Replacement for Universal Visual Recognition. *International Journal of Computer Vision*, 133(2), 456–467.
- [9] Xu, D., et al. (2021). Scene Graph Generation by Iterative Message Passing. *Computer Vision and Image Understanding*, 207, 103196.
- [10] Shi, L., et al. (2022). Skeleton-Based Action Recognition with Multi-Stream CNN. *IEEE Transactions on Neural Networks and Learning Systems*, 33(6), 2367–2379.
- [11] Zhao, B., et al. (2022). Knowledge Graph-Augmented Visual Recognition for Ambiguous Categories. *IEEE Transactions on Image Processing*, 31(7), 4567–4578.
- [12] Chen, X., et al. (2023). Graph Attention Networks for Fine-Grained Visual Recognition. *Computer Vision and Image Understanding*, 234, 103578.
- [13] Wang, X., et al. (2023). Cross-Modal Graph Fusion for Image-Text Matching and Retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4), 3456–3467.
- [14] Pan, Z., et al. (2024). Dynamic Graph Topology Learning for Visual Scene Understanding. *International Journal of Computer Vision*, 132(5), 1234–1245.
- [15] Li, Z., et al. (2025). GNN-Prototype Learning for Few-Shot Visual Recognition. *Computer Vision and Image Understanding*, 259, 103967.
- [16] DeTone, D., et al. (2021). SuperPoint: Self-Supervised Interest Point Detection and Description. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence, 43(8), 2678–2689.
- [17] Ding, M., et al. (2022). Deformable Convolutional Features for Cross-Domain Image Matching. *IEEE Transactions on Image Processing*, 31(3), 1567–1578.
- [18] Noh, H., et al. (2022). DELF: Large-Scale Image Retrieval with Attentive Deep Local Features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6), 3123–3134.
- [19] Radenović, F., et al. (2023). Fine-Tuning CNN Image Retrieval with No Human Annotation. *International Journal of Computer Vision*, 131(2), 234–245.
- [20] Cao, J., et al. (2024). Graph-Based Image Retrieval Re-Ranking. *Computer Vision and Image Understanding*, 248, 103756.
- [21] Wang, F., et al. (2025). Foundation Model Features for Universal Image Retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(2), 1234–1245.
- [22] Dosovitskiy, A., et al. (2021). An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale. *ICLR 2021*.
- [23] He, K., et al. (2022). Masked Autoencoders are Scalable Vision Learners. *Proceedings of CVPR 2022*, 16000–16009.
- [24] Radford, A., et al. (2022). Learning Transferable Visual Models from Natural Language Supervision. *Proceedings of ICML 2021*; Extended *IEEE TPAMI*, 44(8), 3456–3467.
- [25] Chen, T., et al. (2023). Contrastive Learning with Graph Structure for Visual Representation. *Computer Vision and Image Understanding*, 236, 103612.
- [26] Li, Y., et al. (2024). DINOv2: Learning Robust Visual Features Without Supervision. *Transactions on Machine Learning Research*, 2024, 1–35.
- [27] Zhang, Y., et al. (2025). Unified Visual Understanding Framework for Detection, Segmentation, and Recognition. *International Journal of Computer Vision*, 133(4), 678–689.
- [28] Xi, Z., et al. (2023). The Rise and Potential of LLM-Based Agents. *arXiv:2309.07864*; *IEEE TPAMI Extension*, 46(3), 1234–1256.
- [29] Zhu, M., et al. (2023). VPT: Learning to Act by Watching Unlabelled Online Videos. *Advances in NeurIPS*, 35, 14337–14351.
- [30] Wang, X., et al. (2024). VisualAgent: Multi-Agent Visual Reasoning for Structured Image Understanding. *Computer Vision and Image Understanding*, 250, 103823.
- [31] Gupta, T., et al. (2024). VisProg: Compositional Visual Programming for Complex Queries. *Proceedings of CVPR 2024*, 1234–1244.
- [32] Schmidgall, S., et al. (2024). AgentClinic: A Multimodal Agent Benchmark. *NeurIPS 2024*.
- [33] Liu, Y., et al. (2024). Confidence-Based Routing Agent for Visual Recognition. *IEEE Transactions on Image Processing*, 33(5), 2345–2356.
- [34] Ferber, D., et al. (2024). GPT-4V Multi-Agent Visual Analysis for Pathology Interpretation. *Nature Communications*, 15(1), 4743.
- [35] Yang, Z., et al. (2025). Fully Agentic Image Understanding: From Low-Level Features to High-Level Semantics. *Computer Vision and Image Understanding*, 265, 104089.
- [36] M. Sivaram, M. Kaliappan, S. J. Shobana, M. V. Prakash, V. Porkodi, K. Vijayalakshmi, S. Vimal, and A. Suresh, "Secure storage allocation scheme using fuzzy based heuristic algorithm for cloud," *J. Ambient Intell. Humaniz. Comput.*, vol. 12, pp. 5609–5617, 2021.
- [37] E. Mariappan, M. Kaliappan, and S. Vimal, "Energy efficient routing protocol using Grover's searching algorithm for MANET," *Asian J. Inf. Technol.*, vol. 15, pp. 4986–4994, 2016.
- [38] S. Vimal, Y. H. Robinson, M. Kaliappan, A. A. Alaboudi, A. K. Luhach, and N. Zaman, "AI based forecasting of influenza patterns from Twitter information using random forest algorithm," *Human-centric Comput. Inf. Sci.*, vol. 11, no. 33, pp. 1–14, 2021.
- [39] S. Vimal, Y. H. Robinson, M. Kaliappan, K. Vijayalakshmi, and S. Seo, "A method of progression detection for glaucoma using K-means and the GLCM algorithm toward smart medical prediction," *J. Supercomput.*, 2021, doi: 10.1007/s11227-021-03757-w.
- [40] P. Sampath, N. Sasikaladevi, S. Vimal, and M. Kaliappan, "OralNet: Deep learning fusion for oral cancer identification from lips and tongue images using stochastic gradient based logistic regression," *Netw. Model. Anal. Health Inform. Bioinform.*, vol. 13, no. 1, art. 24, 2024.
- [41] G. Sivakumar, M. Kaliappan, and L. J. Julius, "Enhancing the performance of MANET using EESCP," in *Proc. Int. Conf. Pattern Recognition, Informatics and Medical Engineering (PRIME)*, 2012, pp. 221–230.
- [42] S. Vimal, V. Jeyabalaraja, P. Subbulakshmi, A. Suresh, M. Kaliappan, and S. Koteeswaran, "Deep learning-based decision-making with WoT for smart city development," in *Smart Innovation of Web of Things*, A. Jain, R. G. Crespo, and M. Khari, Eds. Boca Raton, FL, USA: CRC Press, 2020, pp. 51–62.
- [43] M. Kaliappan and B. Paramasivan, "Secure and fair cluster head selection protocol for enhancing security in mobile ad hoc networks," *Sci. World J.*, vol. 2014, art. 608984, 2014.
- [44] S. Vimal, M. Kaliappan, A. Suresh, P. Subbulakshmi, S. Kumar, and D. Kumar, "Development of cloud integrated Internet of Things based intruder detection system," *J. Comput. Theor. Nanosci.*, vol. 15, no. 11–12, pp. 3565–3570, 2018.
- [45] C. Karpagavalli and M. Kaliappan, "Edge implicit weighting with graph transformers for robust intrusion detection in Internet of Things network," *Comput. Security*, vol. 150, 2025.
- [46] S. Gopikumar, J. R. Banu, Y. H. Robinson, S. Raja, S. Vimal, D. Pelusi, and M. Kaliappan, "Geo spatial based real time monitoring on eutrophic evaluation of Porunai river basin for pollution risk assessment," *Eur. J. Remote Sens.*, pp. 1–13, 2022, doi: 10.1080/22797254.2021.2025152.
- [47] M. Kaliappan, E. Mariappan, and B. Revathi, "Unified vision transformer for multimodal medical image

- classification," *Biomed. Signal Process. Control*, art. 110431, 2026.
- [48] M. Kaliappan, E. Mariappan, M. V. Prakash, and B. Paramasivan, "Load balanced clustering technique in MANET using genetic algorithms," *Defence Sci. J.*, vol. 66, no. 3, pp. 251–258, 2016