ENCODING ROBUST TOPOLOGICAL SIGNATURES FOR HYPERDIMENSIONAL COMPUTING

Anonymous authors
Paper under double-blind review

ABSTRACT

Hyperdimensional computing (HDC) enables lightweight prototype-based classification but pixel-based encoders are brittle under geometric and pixel-level corruption. We introduce a topology-guided HDC classifier that combines local HOG features, rotation/translation/scale-stable spatial-pyramid Zernike descriptors, and a robust hole-count compatibility score. Hole candidates are generated from multiscale bounded background components and validated using a spoke-wheel enclosure test. Their counts are converted into class-conditional topology scores learned from clean training data and fused with cosine similarities from HOG and Zernike HDC prototypes. On MNIST and EMNIST Letters, the method substantially improves robustness over naive HDC and often outperforms a compact CNN under Gaussian noise, salt-and-pepper noise, and severe scaling, without corruption-specific training. OnlineHD updates further improve clean and occlusion performance, although adaptation can reduce robustness under some noise settings. These results show that explicit, interpretable topology provides a practical complement to appearance-based HDC.

1 INTRODUCTION

Hyperdimensional computing (HDC) represents data with high-dimensional distributed vectors and performs classification using lightweight prototype-based operations (Kanerva, 2009; Rahimi et al., 2016; Kleyko et al., 2023). Its simple arithmetic, modest memory requirements, and support for online updates make HDC attractive for resource-constrained learning. These benefits, however, depend strongly on the input encoding. Standard image encoders bind pixel intensities to position hypervectors and bundle them into a single representation (Smets et al., 2024). Because small geometric or photometric changes can alter many pixel–position associations simultaneously, such encoders can be brittle under distribution shift. In our experiments, for example, naive HDC accuracy on MNIST falls from 95.5% on clean images to 9.8% under Gaussian noise with $\sigma = 0.1$; complete confusion matrices are provided in Appendix 6.

We investigate whether explicit shape and topology can provide more stable evidence for HDC image classification. Topological properties such as connectedness and hole count are preserved under homeomorphisms, including rotation, translation, and isotropic scaling. They therefore offer an interpretable complement to appearance features when pixel statistics are disrupted. Topological structure has previously been studied both as an explicit image representation (Carlsson et al., 2008; Carlsson, 2020; Zhou et al., 2017) and as an emergent property of neural representations (Naitzat et al., 2020). Nevertheless, explicit topological evidence has received comparatively little attention in HDC, where prior image encoders primarily represent pixel appearance or implicit global structure (Smets et al., 2024; Vergés et al., 2025).

We propose a topology-guided HDC classifier that combines three complementary sources of evidence. First, HOG captures local stroke and gradient structure. Second, spatial-pyramid Zernike magnitudes summarize global and coarse regional shape after centering, scale normalization, and canonical orientation. Third, we introduce a robust hole detector that generates bounded-background candidates across multiple morphological scales and validates their local enclosure using a spoke-wheel test.

The detected hole count is mapped to a topology category $z(x) \in 0, 1, 2+$, and clean training data are used to estimate a class-conditional compatibility score $logP(z(x)|c)$. This formulation uses topology for the question it answers most directly—whether the observed structural category

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

is compatible with a class—rather than requiring sparse hole features to distinguish every class independently. HOG and Zernike descriptors are mapped to bipolar hypervectors and scored against class prototypes; their cosine similarities are then fused with the topology-compatibility score. The Zernike weight is selected on held-out validation data, while the topology weight is fixed across the pre- and post-OnlineHD stages.

Our contributions are:

- We introduce an HDC image classifier that combines local appearance, canonicalized outer-shape information, and an explicit, class-conditional topology-compatibility score.
- We develop a robust hole-extraction procedure based on multiscale connected-component proposals, spoke-wheel enclosure validation, and persistence across morphological scales.
- We evaluate the method on MNIST and EMNIST Letters under rotation, Gaussian noise, salt-and-pepper noise, cutout, and zoom. The proposed approach substantially improves robustness over naive HDC and often outperforms a compact CNN under pixel-level corruption, without corruption-specific training.

2 TOPOLOGICAL SIGNATURE ENCODING

We cast the problem of image classification as a shape identification problem - *given different observations of an object under nuisance transformation and noise, our HD classifier has to decide whether they represent the same underlying shape.*

For a grayscale image $x : \Omega \to \mathbb{R}$, the true shape is a compact subset $\mathcal{A} \subset \mathbb{R}^2$ of the foreground region. An observed image produces an estimate $\hat{\mathcal{A}}$ via thresholding/cleanup. Shape identification typically considers the equivalence class

$$[\mathcal{A}] = g(\mathcal{A}) : g \in G \tag{1}$$

where G is the nuisance group. Our goal is to construct an ideal map

$$\Psi(\mathcal{A}) \in \{-1, +1\}^D \tag{2}$$

such that $\Psi(g(\hat{\mathcal{A}})) \approx \Psi(\hat{\mathcal{A}})$ for all $g \in G$ and Ψ is discriminative across different classes of shapes.

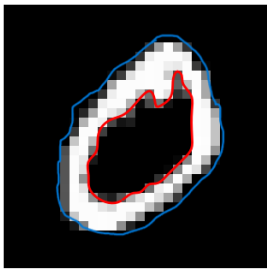


Figure 1: Outer contour (blue) and hole contour (red).

The shape $\mathcal{A}$ in our case could be framed in two parts - global geometry through the shape of outer boundary $\mathcal{A}_O$ and geometry of the holes (if any) present inside the shape, $\mathcal{A}_{\mathcal{H}_i}$, where $\mathcal{H}_i$ represents the i^{th} hole, $0 \leq \mathcal{H}_i \leq \mathcal{H}_{max}$. The outer contour is the usual object silhouette information which represents the boundary/appearance-based invariants commonly used in classical shape matching.

Holes, on the other hand, have been utilized as primary topological features (Betti-1 in the binary region) due to their invariance under any affine transformation (Edelsbrunner and Harer, 2008). However, the location and shape of the individual hole is dependent on the outer contour, i.e. the scale of the overall shape dictates the relative location of the hole. Therefore, the first step is to find the existence of holes via connected component analysis. Given that some images do not contain holes naturally (for e.g. the mnist class 1), we wish to infer the topology category from the inferred holes in a given image (Fig. 1 shows the outer and hole contours). The methodology is described in detail below:

2.1 METHOD OVERVIEW

Given an image x, we compute HOG and spatial-pyramid Zernike descriptors and map each to a bipolar hypervector using randomized projection. Their class scores are cosine similarities to channel-specific HDC prototypes. Separately, a multiscale spoke-validated extractor estimates the number of foreground holes and produces a class-conditional topology-compatibility score. The three scores are combined by late fusion; HDC prototypes may subsequently be refined using OnlineHD updates.

2.2 APPEARANCE AND OUTER-SHAPE CHANNELS

We use HOG to represent local stroke structure and spatial-pyramid Zernike (SPZ) magnitudes to represent global and coarse regional shape. Images are centered, scale-normalized, and rotated to the canonical outer-contour frame. The spatial pyramid version of the Zernike descriptor retains Zernike moments' strong global invariance properties while recovering spatial layout information by computing moments not only on the full image but also on a coarse grid of subregions. We compute Zernike magnitudes on the full image and on a 2×2 grid and concatenate the results:

$$\phi_{\mathrm{SPZ}}(x) = [\phi_0(x); \phi_1(x); \ldots; \phi_g(x)]. \tag{3}$$

Both HOG and SPZ vectors are independently mapped to bipolar hypervectors by randomized projection. Implementation details and descriptor hyperparameters are provided in Appendix 6.

2.3 ROBUST HOLE EXTRACTION AND TOPOLOGY COMPATIBILITY

Let $A \subset \Omega$ denote the normalized binary foreground mask. In the uncorrupted setting, holes are bounded connected components of the background $\Omega \setminus A$, i.e., components that do not intersect the image boundary. Because noise and discretization may introduce small gaps in an otherwise enclosed region, we generate candidates at multiple morphological scales and validate them using a spoke-wheel enclosure test.

Multiscale hole proposals For each closing radius $r \in \mathcal{R}$, let

$$A_r = \mathrm{Close}_r(A) \tag{4}$$

denote the foreground after binary closing, with $A_0 = A$. Candidate holes at scale r are defined as

$$\mathcal{B}_r = \{H \in \mathrm{CC}(\Omega \setminus A_r) : H \cap \partial\Omega = \emptyset,\ a_{\min} \le |H| \le a_{\max}\}, \tag{5}$$

where CC denotes connected components and the area constraints remove small artifacts and implausibly large background regions.

Spoke-wheel enclosure validation For each candidate H, we select an interior seed q_H as the maximizer of the distance transform within H. From q_H, we cast N_θ uniformly spaced spokes with angles

$$\theta_j = \frac{2\pi(j-1)}{N_\theta}, \qquad j = 1, \ldots, N_\theta. \tag{6}$$

Let d_j^H denote the distance at which spoke j exits the candidate region, and let d_j^A denote the first intersection with the original foreground mask A. We define the local boundary-support indicator

$$b_j(H) = \mathbf{1}\left\{\left|d_j^A - d_j^H\right| \le \tau\right\}, \tag{7}$$

where τ permits small boundary gaps. The enclosure score is

$$e(H) = \frac{1}{N_\theta} \sum_{j=1}^{N_\theta} b_j(H). \tag{8}$$

We additionally compute $g(H)$, the largest circular run of unsupported spokes divided by N_θ. Candidates are retained only when

$$e(H) \ge \eta_e \qquad \text{and} \qquad g(H) \le \eta_g. \tag{9}$$

Unlike testing whether a spoke intersects foreground anywhere along its path, this local support criterion prevents distant, unrelated strokes from incorrectly validating an open concavity.

Multiscale persistence Candidates detected at different closing radii are matched using overlap and centroid proximity. For a matched candidate group, we define its persistence as

$$\rho(H) = \frac{|\{r \in \mathcal{R} : H \text{ is detected at scale } r\}|}{|\mathcal{R}|}. \tag{10}$$

The final hole set $\widehat{\mathcal{H}}(A)$ contains candidates satisfying the enclosure, angular-gap, and persistence criteria. This hybrid construction uses connected components to preserve the topological definition of a hole, while spoke validation and multiscale persistence improve robustness to small gaps and pixel-level artifacts.

Topology compatibility Rather than requiring the sparse hole representation to distinguish every class directly, we map the detected hole count to the topology category

$$z(x) = \min\left(\left|\widehat{\mathcal{H}}(A_x)\right|, 2\right) \in \{0, 1, 2+\}. \tag{11}$$

Using only the clean training set, we estimate the class-conditional topology distribution with Laplace smoothing:

$$\widehat{P}(z \mid c) = \frac{N_{c,z} + \lambda}{\sum_{z'} N_{c,z'} + K\lambda}, \tag{12}$$

where $N_{c,z}$ is the number of training samples from class c assigned to topology category z and $K = 3$ are the number of topology categories. The resulting compatibility score is

$$s_c^{\text{topo}}(x) = \log \widehat{P}(z(x) \mid c)\,. \tag{13}$$

2.4 Score fusion and reliability weighting

The HOG and Zernike channels capture complementary appearance and shape information. HOG emphasizes local stroke and gradient structure, whereas Zernike moments summarize global and coarse regional geometry. The robust hole extractor provides a separate topology-compatibility score based on the detected hole-count category. Because the relative contribution of the Zernike channel is not known a priori, we select its weight using a held-out validation set.

Let x denote an input image and let $c \in \{1, \ldots, C\}$ index the class. We construct two bipolar hypervector representations:

$$\mathbf{HV}^{\text{hog}}(x), \mathbf{HV}^{\text{zer}}(x) \in \{-1, +1\}^D, \tag{14}$$

corresponding to the HOG and Zernike descriptors, respectively. For each descriptor type $t \in \{\text{hog}, \text{zer}\}$ and class c, we maintain a class prototype

$$\mathbf{P}_c^t \in \mathbb{Z}^D, \tag{15}$$

initially constructed by bundling the training hypervectors and optionally refined using OnlineHD updates.

The class score for descriptor type t is the cosine similarity between the sample hypervector and its class prototype:

$$s_c^t(x) = \frac{\left\langle \mathbf{HV}^t(x), \mathbf{P}_c^t \right\rangle}{\left\|\mathbf{HV}^t(x)\right\|_2 \left\|\mathbf{P}_c^t\right\|_2 + \varepsilon}, \tag{16}$$

where $\varepsilon > 0$ provides numerical stability.

The topology channel does not construct a separate class prototype. Instead, the detected hole count is mapped to a topology category $z(x) \in \{0, 1, 2+\}$, and its class-compatibility score is

$$s_c^{\text{topo}}(x) = \log \widehat{P}(z(x) \mid c)\,, \tag{17}$$

where $\widehat{P}(z \mid c)$ is estimated once from the clean training set as described in Section 2.3.

We combine the three sources of evidence using

$$S_c(x; \alpha, \beta) = s_c^{\text{hog}}(x) + \alpha s_c^{\text{zer}}(x) + \beta s_c^{\text{topo}}(x), \tag{18}$$

where $\alpha \geq 0$ controls the contribution of the Zernike channel and $\beta \geq 0$ is a fixed weight for topology compatibility. We select α on the held-out validation set:

$$\alpha^* = \arg\max_{\alpha \in \mathcal{A}} \frac{1}{|\mathcal{V}|} \sum_{(x,y) \in \mathcal{V}} \mathbf{1}\left[y = \arg\max_c S_c(x; \alpha, \beta)\right], \tag{19}$$

where $\mathcal{V}$ is the validation set and $\mathcal{A}$ is a predefined search grid. The topology weight β is held fixed across the pre- and post-OnlineHD stages because the topology-compatibility model is estimated once and is not modified during prototype adaptation. After OnlineHD updates, only α is reselected using the same validation protocol. The final prediction is

$$\widehat{y}(x) = \arg\max_{c \in \{1, \ldots, C\}} S_c(x; \alpha^*, \beta). \tag{20}$$

2.5 OPERATIONAL COMPLEXITY

Let

- $P = HW$: pixels in the normalized image;
- B: number of HOG orientation bins;
- M: number of retained Zernike moments;
- S: spatial pyramids
- G: number of Zernike regions, including the global region;
- R: number of morphological closing scales;
- K: number of candidate holes;
- N_θ: number of spokes per candidate; and
- L: average number of samples along each spoke.

HOG consists of gradient computation, orientation voting and block normalization, where for every pixel P, the gradient magnitude and orientation is computed and projected into the bins, accumulating $\mathcal{O}(BP)$ operations. Standard computation of a single Zernike moment over an image region takes $\mathcal{O}(P)$ time. For M moments repeated over G regions, it is $\mathcal{O}(GMP)$. Since we employ a spatial pyramid, the Zernike computation has $\mathcal{O}(SGMP)$ operations. The morphological closing across R different scales provides $\mathcal{O}(RP)$ operations. Finding K candidate holes over P pixels leads to $\mathcal{O}(KP)$ operations. Finally, sampling along N_{theta} radial spokes along L samples for K candidate holes leads to $\mathcal{O}(KN_\theta L)$ operations. Therefore, the overall complexity of the feature extraction can be given as $\mathcal{O}(BP + SGMP + RP + KP + KN_\theta L)$.

2.6 THEORETICAL UNDERPINNINGS

Theorem 2.1. *The spatial pyramid outer shape descriptor is invariant to rotation, translation and scale.*

Proof. Assume that under any similarity transform $tr(p) = aQp+t$ with $a > 0$, $Q \in SO(2)$, and $t \in \mathbb{R}^2$, the pyramid regions transform equivariantly (i.e., $tr(Re_i) = Re_i$ after normalization/resampling), and that $\mathcal{N}_{Re}$ is computed from f on Re so that $c_{tr(Re)} = tr(c_{Re})$ and $s_{tr(Re)} = a\, s_{Re}$. Then for the transformed image $f^t r(p) = f(tr^{-1}p)$,

Translation invariance follows because centering by c_{Re} satisfies $c_{g(Re)} = aQc_{Re} + t$, hence $(p-c_{g(Re)})/s_{g(Re)} = Q(p'-c_{Re})/s_{Re}$ for $p' = g^{-1}p$. Scale invariance follows from $s_{g(Re)} = as_{Re}$, so normalization cancels the factor a. After centering and scaling, a rotation Q acts as a rotation on the unit disk coordinates: $u \mapsto Qu$. Zernike moments satisfy the standard rotation law

where Q is rotation by angle α, hence $|Z_{n,m}|$ is rotation invariant. Applying the same argument region-wise and concatenating over the pyramid yields the result.

□

Theorem 2.2. *Let $A \subset \Omega$ be the ideal foreground and $\tilde{A}$ be the corrupted estimate. Let A_r denote the binary closing with r radius $A_r = Close_r(\tilde{A})$. For a bounded background candidate $H \subset \Omega \setminus A_r$, let q_H be an interior seed. For direction θ, denote the distance from q_H to the candidate boundary $d_H(\theta)$ and the distance to the first intersection with the original estimated foreground $\tilde{A}$ is given as $d_A(\theta)$. Given the true hole H^* satisfies the following assumptions:*

- *its boundary gaps in $\tilde{A}$ are repaired by closing at some $r \in [0, R_{max}]$;*
- *closing at radius r does not merge H^* with another background or image exterior;*
- *boundary localization error satisfies $|d_A(\theta) - d_H(\theta)| \leq \tau$ except on an angular set of relative measure at most ϵ;*
- *the largest unsupported angular interval has relative length at most γ.*

Then up to angular discretization error,

$$e_N(H^*) \geq 1 - \epsilon - \frac{2}{N_\theta}, g_N(H^*) \leq \gamma + \frac{2}{N_\theta}. \tag{21}$$

Thus, H^* is accepted whenever,

$$\eta_e \leq 1 - \epsilon - \frac{2}{N_\theta}, \eta_g \geq \gamma + \frac{2}{N_\theta}. \tag{22}$$

Proof. The proof is provided in Appendix 6.1. □

3 RESULTS

We evaluate the proposed topological encoding method on two standard vision benchmarks spanning increasing visual complexity: MNIST (LeCun et al., 2002) and EMNIST (Cohen et al., 2017); on an Ubuntu GPU server machine with two NVIDIA A6000 cards and Intel® Xeon(R) Bronze 3206R CPU @ 1.90GHz × 16. MNIST provides a controlled setting of centered grayscale digits, EMNIST extends this regime to a larger and more diverse set of handwritten characters. Specifically, we utilize the letters of the EMNIST dataset, which have small strokes that vanish at low resolution. Together, these datasets allow us to assess both in-distribution accuracy and robustness as the input distribution shifts from clean binary-like shapes to natural scenes. The algorithm pseudocode is provided in Appendix 6.1 and the hyperparameter settings are provided in Appendix 6.1.

We further stress-test generalization by evaluating under common distribution shifts induced by synthetic corruptions applied to the test sets. This provides a realistic measure of generalization of the method to unseen variations. Specifically, we generate corrupted variants of each test dataset using: (i) in-plane rotations up to a specified maximum angle, (ii) additive Gaussian noise with controlled standard deviation, (iii) salt-and-pepper impulse noise at a fixed flip probability, (iv) cut-out occlusions that mask a contiguous square region of fixed size, and (v) zoom transformations that rescale the foreground content and re-center it on the original canvas. These perturbations target complementary failure modes—geometric misalignment, stochastic pixel noise, sparse outliers, partial occlusion, and scale changes—allowing us to quantify how well the proposed encoding preserves class separability when the input deviates from the training distribution.

3.1 MNIST DATASET

We start by analyzing the learned topology distribution of the MNIST dataset. The robust extractor detects no holes in 62.0% of the MNIST training images and identifies 0.44 holes per image on average, with at most four candidates retained. The learned class-conditional distributions broadly agree with the expected topology of the digits. Classes 0, 6, and 9 are assigned predominantly to the one-hole category, with probabilities of 87.8%, 76.4%, and 75.8%, respectively. Class 8 is assigned predominantly to the two-or-more-hole category, with a probability of 62.8%. Classes 1, 3, 4, 5, and 7 are predominantly hole-free, with probabilities ranging from 93.9% to 100%. Class 2 exhibits greater writing-style and extraction variability: 63.0% of its samples are categorized as hole-free and 36.6% as containing one hole. More generally, the topology distributions show that hole count is informative but not deterministic. Atypical categories can arise from handwriting variation, imperfect binarization, or spurious and missed holes. This variability motivates using the learned class-conditional compatibility score rather than imposing a hard mapping between digit classes and hole counts.

The topology distributions are estimated from clean training data with Laplace smoothing and are not manually specified. At inference, the observed hole-count category provides soft class evidence rather than excluding classes whose topology differs from their canonical form. This allows the topology term to complement the HOG and Zernike channels while remaining tolerant of extraction uncertainty.

Table 1 compares naive HDC, Topology-guided HDC, and a compact five-layer CNN. All models are trained using clean MNIST only and evaluated using test data under specific corruptions. Topology-guided HDC combines HOG and canonicalized spatial-pyramid Zernike evidence with a class-conditional topology score derived from the detected hole-count category. It substantially outperforms

Table 1: MNIST test accuracy (%) under clean and corrupted settings. All models use clean training data only; no corruption-specific augmentation is applied. "Initial" denotes bundled HDC prototypes, and "OnlineHD" denotes the prototypes after 20 update epochs.

	Naive HDC		Topology-guided HDC		CNN
Setting	**Initial**	**OnlineHD**	**Initial**	**OnlineHD**	**Clean-trained**
Clean	82.10	95.50	93.96	97.68	**99.10**
Rotation (20°)	20.50	18.30	93.89	97.61	**99.16**
Gaussian noise ($\sigma = 0.1$)	6.86	9.78	91.53	**93.80**	11.35
Gaussian noise ($\sigma = 0.2$)	7.53	6.91	90.15	**92.80**	11.01
Salt-and-pepper ($p = 0.1$)	43.86	65.84	91.15	**92.40**	40.10
Cutout (size $= 12$)	62.19	44.47	93.22	97.03	**98.52**
Zoom (scale $= 0.75$)	19.86	14.89	90.74	**93.71**	92.20
Zoom (scale $= 0.5$)	15.43	13.02	84.15	**87.19**	52.85

Table 2: Hole-channel-only test accuracy (%) before OnlineHD adaptation. Bold values indicate the best result for each setting.

Ablation	**Clean**	**Rotation**	**Gaussian**	**Salt-and-pepper**	**Cutout**	**Zoom**
A0: Original	28.50	28.64	28.25	23.91	18.99	28.42
A1: Multiscale	35.73	35.84	33.51	31.97	22.18	35.05
A2: Multiscale + spoke	36.62	36.52	33.85	32.72	22.51	35.70
A3: Multiscale + spoke + pyramid	**40.41**	**40.47**	**36.22**	**37.02**	**24.61**	**38.22**
A4: Exact + spoke	28.50	28.64	28.25	23.87	19.58	28.42

naive HDC under every evaluated corruption, both with initial bundled prototypes and after OnlineHD adaptation.

The CNN achieves the highest clean and rotation accuracy and is marginally better under the small cutout. In contrast, Topology-guided HDC is markedly more robust to pixel noise: after OnlineHD, it obtains 93.80% and 92.80% under Gaussian noise, compared with 11.35% and 11.01% for the CNN, and 92.40% under salt-and-pepper noise, compared with 40.10%. It also outperforms the CNN under both zoom settings, with the largest margin at scale 0.5 (87.19% versus 52.85%).

Overall, the result reveals that incorporating topology-guided shape information into an HDC framework yields substantial robustness over a naive HDC baseline across diverse corruptions, even before online adaptation and consistently after online training. Second, while the CNN achieves superior accuracy on clean MNIST, Topology-guided HDC can offer markedly better robustness for several common corruption types—especially those that disrupt local pixel statistics—without requiring corruption-specific training. These results underscore the value of topology-guided representations as a complementary route to robustness, particularly in settings where retraining or extensive augmentation is undesirable.

3.1.1 Ablation Study

We ablate the components of the hole-extraction pipeline on MNIST while keeping the HOG and Zernike encoders, HDC dimension, data split, and OnlineHD settings fixed. We consider five variants: A0 (original) detects holes as bounded background components without additional validation; A1 (multiscale) generates candidates after morphological closing at multiple scales and retains persistent detections; A2 (spoke) adds spoke-wheel enclosure validation to the multiscale candidates; A3 (spoke+pyramid) additionally encodes the canonical spatial location of each hole using a spatial pyramid; and A4 (exact+spoke) applies spoke validation to the original, single-scale connected-component proposals. Results are reported before (in Table 2) and after (in Table 3) 20 epochs of OnlineHD adaptation under clean, rotation, Gaussian-noise, salt-and-pepper, cutout, and zoom conditions.

The ablations produce nearly identical clean and rotation accuracy, indicating that the proposed hole-processing changes do not compromise in-distribution performance or the existing geometric invariance. Before OnlineHD, no variant dominates all corruptions. A0 performs best under Gaussian

Table 3: Hole-channel-only test accuracy (%) after 20 epochs of OnlineHD adaptation. Bold values indicate the best result for each setting.

Ablation	Clean	Rotation	Gaussian	Salt-and-pepper	Cutout	Zoom
A0: Original	38.92	37.39	34.45	27.75	18.02	35.79
A1: Multiscale	41.87	**43.12**	35.37	**38.16**	21.08	**38.93**
A2: Multiscale + spoke	38.50	39.70	34.81	36.04	21.94	32.06
A3: Multiscale + spoke + pyramid	**47.21**	42.53	30.78	38.03	**27.22**	31.20
A4: Exact + spoke	39.37	38.93	**35.53**	26.54	17.71	38.06

noise (92.51%), A1 under salt-and-pepper noise (91.53%), and A4 under cutout (84.50%); A0 and A4 tie under zoom (90.74%). After OnlineHD, A4 achieves the best Gaussian-noise and cutout accuracy (94.61% and 89.04%), A1 performs best under salt-and-pepper noise (93.80%), and A3 is marginally best under zoom (93.71%). Averaged across the five corrupted settings after adaptation, the variants differ by less than 0.3 percentage points, with A1 obtaining the highest mean accuracy (93.38%). Thus, multiscale and spoke-based extraction provide corruption-specific gains but do not yield a uniformly superior fused classifier.

The hole-only results more clearly expose the effect of the individual components. Before OnlineHD, A3 achieves the highest hole-channel accuracy under every evaluated condition, improving clean accuracy from 28.50% for A0 to 40.41% and improving Gaussian-noise accuracy from 28.25% to 36.22%. This indicates that multiscale proposals, spoke validation, and spatial information produce a more discriminative hole representation. After adaptation, however, the best variant depends on the corruption: A3 remains strongest on clean data and cutout, A1 performs best under rotation, salt-and-pepper noise, and zoom, and A4 performs best under Gaussian noise. These results suggest that the detailed hole representation is informative in isolation but partly redundant with the stronger HOG and Zernike channels in the fused classifier. OnlineHD also appears to change the relative value of the topology cues, motivating the use of a fixed topology-compatibility term rather than independently adapting a sparse hole prototype.

Overall, the ablation supports two conclusions. First, multiscale proposals substantially improve the informativeness of the extracted hole structure, while spoke validation provides additional benefits in selected noisy conditions. Second, adding every component does not consistently improve end-to-end accuracy: the full A3 representation improves the hole channel but is not the strongest fused variant after adaptation. We therefore use the robust extraction components to estimate hole-count compatibility and treat the finer spatial-pyramid representation as an auxiliary ablation rather than claiming universal improvement.

3.2 EMNIST DATASET

Table 4: Test accuracy for Topology-guided HDC and CNN for EMNIST letters dataset before and after online training for 20 epochs under clean and corrupted settings.

	Topology-guided HDC		CNN
Setting	**Before train**	**After train**	**After train**
Baseline (clean)	76.12	89.71	**93.21**
Rotation (20°)	76.09	89.91	**93.21**
Gaussian noise ($\sigma = 0.1$)	71.5	**72.69**	3.84
Salt-and-pepper ($p = 0.1$)	5.01	4.06	**31.51**
Cut out ($size = 4$)	74.22	87.22	**91.01**
Cut out ($size = 8$)	70.06	**82.34**	82.11
Cut out ($size = 12$)	62.98	**70.09**	67.98
Zoom (scale = 0.75)	4.51	**5.12**	3.85
Zoom (scale = 0.5)	4.31	**5.08**	3.84

For EMNIST Letters, naive HDC performs substantially worse than on MNIST and is therefore omitted from the main comparison. As shown in Table 4, the CNN achieves the highest accuracy on

clean and rotated inputs (93.21%), while Topology-guided HDC is markedly more robust to Gaussian noise, attaining 71.50% initially and 72.69% after OnlineHD, compared with 3.84% for the CNN. OnlineHD also improves clean accuracy from 76.12% to 89.71% and provides competitive or superior performance under larger cutouts, reaching 82.34% at size 8 and 70.09% at size 12. However, this robustness does not extend uniformly across corruptions: performance under salt-and-pepper noise falls to 5.01% initially and 4.06% after adaptation, while all methods remain near chance under zoom. These results indicate that the proposed structural representation is particularly effective under Gaussian noise and substantial occlusion, but remains sensitive to sparse impulse noise and severe rescaling on the more heterogeneous EMNIST Letters dataset.

4 Discussion

This work introduces a topology-guided HDC classifier that combines two prototype-based hypervector channels with an explicit topology-compatibility score. HOG captures local stroke information, while canonicalized spatial-pyramid Zernike descriptors represent global and coarse regional shape. Holes are detected using multiscale connected-component proposals, spoke-wheel enclosure validation, and persistence across morphological scales. Rather than encoding individual holes as hypervectors, the detected count is mapped to a topology category—zero, one, or two or more holes—and converted into a class-conditional compatibility score learned from clean training data. This score is fused with the HOG and Zernike cosine similarities using a fixed topology weight and a validation-selected Zernike weight.

Across MNIST and EMNIST Letters, the proposed classifier is most effective when corruptions disrupt local pixel statistics while preserving recoverable coarse structure. On MNIST, it substantially outperforms naive pixel-based HDC across all evaluated shifts and is markedly more robust than the compact CNN under Gaussian and salt-and-pepper noise and severe zoom. On EMNIST, it remains particularly effective under Gaussian noise and larger cutouts, although it performs poorly under salt-and-pepper noise and severe zoom. These results suggest that canonicalized shape and hole-count compatibility provide useful complementary evidence, but their effectiveness depends on whether preprocessing can recover stable foreground structure after corruption.

OnlineHD generally improves clean, rotation, and occlusion performance by adapting the HOG and Zernike class prototypes. The topology model itself remains fixed because it is estimated once from clean training data. This separation preserves an interpretable structural prior during adaptation, although the EMNIST results show that prototype updates do not uniformly improve robustness across corruption types. Overall, the findings support explicit topology as a lightweight complement to appearance-based HDC rather than as a standalone representation.

4.1 Limitations

However, the results also highlight failure modes. The method depends on reliable foreground normalization and hole extraction. Severe noise, sparse impulse corruption, low contrast, or rescaling can fragment boundaries, close open concavities, or create spurious holes, leading to an incorrect topology category and compatibility score; this limitation is evident for salt-and-pepper noise and severe zoom on EMNIST Letters. The geometric invariance of the Zernike channel is also conditional on stable centering, scaling, and orientation estimation and does not cover nonrigid deformation, perspective change, blur, compression, or substantial background clutter. Performance may be sensitive to thresholding and cleanup parameters, morphological scales, spoke-enclosure thresholds, Zernike order, spatial-pyramid resolution, and the fixed topology weight; these choices are not exhaustively characterized. Moreover, the topology model represents only the coarse categories of zero, one, or two or more holes and therefore cannot distinguish classes with similar hole-count distributions without appearance and shape evidence. Finally, although HDC prototype scoring is lightweight, image normalization, Zernike computation, and multiscale hole detection may dominate runtime at higher resolutions. Thus, robustness ultimately depends on whether stable geometric and topological structure remains recoverable after corruption.

5 AI USE STATEMENT

In this work, we used generative AI tools for polishing the written text. We have also used generative AI for coding and checking proofs of theorems. Everything generated by AI has been thoroughly reviewed by authors.

6 REPRODUCIBILITY STATEMENT

We have provided the code as a zipped file along with the submission. The code can be independently run with benchmark datasets.

REFERENCES

Carlsson, G. (2020). Topological methods for data modelling. *Nature Reviews Physics*, 2(12):697–708.

Carlsson, G., Ishkhanov, T., De Silva, V., and Zomorodian, A. (2008). On the local behavior of spaces of natural images. *International journal of computer vision*, 76(1):1–12.

Cohen, G., Afshar, S., Tapson, J., and Van Schaik, A. (2017). Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE.

Edelsbrunner, H. and Harer, J. (2008). Persistent homology -— a survey. In *Surveys on discrete and computational geometry*, volume 453, page 257. Amer Mathematical Society.

Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive computation*, 1(2):139–159.

Kleyko, D., Rachkovskij, D., Osipov, E., and Rahimi, A. (2023). A survey on hyperdimensional computing aka vector symbolic architectures, part ii: Applications, cognitive models, and challenges. *ACM Computing Surveys*, 55(9):1–52.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (2002). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.

Naitzat, G., Zhitnikov, A., and Lim, L.-H. (2020). Topology of deep neural networks. *Journal of Machine Learning Research*, 21(184):1–40.

Rahimi, A., Kanerva, P., and Rabaey, J. M. (2016). A robust and energy-efficient classifier using brain-inspired hyperdimensional computing. In *Proceedings of the 2016 international symposium on low power electronics and design*, pages 64–69.

Smets, L., Van Leekwijck, W., Tsang, I. J., and Latré, S. (2024). An encoding framework for binarized images using hyperdimensional computing. *Frontiers in big data*, 7:1371518.

Vergés, P., Nicolau, A., and Givargis, T. (2025). Learning encoding phasors with fractional power encoding. In *2025 11th International Conference on Computing and Artificial Intelligence (ICCAI)*, pages 537–543. IEEE.

Zhou, Z., Huang, Y., Wang, L., and Tan, T. (2017). Exploring generalized shape analysis by topological representations. *Pattern Recognition Letters*, 87:177–185.

APPENDIX

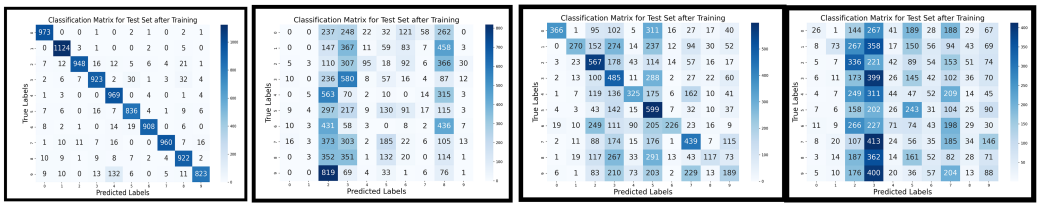


Figure 2: Confusion matrices of MNIST dataset classification after training for 20 epochs using naive HDC implementation with the following configurations: (a) Clean test dataset with accuracy of 0.95; (b) adding gaussian noise of $\sigma = 0.1$ results in accuracy of 0.09; (c) salt-and-pepper noise for 10% of pixels with accuracy of 0.65; and (d) zoom of 0.75 with accuracy of 0.15. The confusion matrices clearly show mode collapse occurring due to added noise.

OUTER SHAPE EXTRACTION AND MOMENT DESCRIPTORS

For the grayscale image x, we isolate the foreground pixels using a global or adaptive threshold map $T : \Omega \to \mathbb{R}$ to form a binary mask $b(p), p \in \Omega$, whereby $\mathcal{C}(b)$ provide the set of connected components of $\{p \in \Omega : b(p) = 1\}$. The shape is then given by the largest connected component

$$\mathcal{A} = \arg\max_{\mathcal{C} \in \mathcal{C}(b)} |C|, b_{\mathcal{A}}(p) = \mathbb{1}\{p \in \mathcal{A}\}. \tag{23}$$

A continuous region is then obtained via pixel embedding,

$$\tilde{Re} = \bigcup_{(i,j)\in\mathcal{A}} \left(\left[i - \frac{1}{2}, i + \frac{1}{2} \right] \times \left[j - \frac{1}{2}, j + \frac{1}{2} \right] \right) \subset \mathbb{R}^2. \tag{24}$$

The outer contour is then defined as the outer boundary component of $\partial\tilde{Re}$. Let $\partial\Omega$ denote the set of border pixels. Let $\mathcal{C}(\bar{b})$ be the set of connected components of $\{p : \bar{b}(p) = 1\}$, and define

$$\mathcal{H}(\mathcal{A}) = \left\{ \mathcal{H} \in \mathcal{C}(\bar{b}) \ : \ \mathcal{H} \cap \partial\Omega = \emptyset \right\}. \tag{25}$$

Each $\mathcal{H} \in \mathcal{H}(\mathcal{A})$ is a hole region. For each hole region $\mathcal{H}$, we extract its boundary (contour) as an isocontour at level $1/2$ of the indicator image of $\mathcal{H}$ (marching squares), yielding an ordered set of points

$$C_{\mathcal{H}} = \{c_1, \ldots, c_{n_{\mathcal{H}}}\} \subset \mathbb{R}^2, \tag{26}$$

which we treat as a closed polygonal curve (with $c_{n_{\mathcal{H}}+1} = c_1$).

Let $C_O = \{o_1, \ldots, o_{n_0}\} \subset \mathbb{R}^2$ denote the outer contour of the foreground (e.g., the longest extracted contour). We define a canonical RTS frame using the outer contour. First the outer centroid $c = \frac{1}{n_0} \sum_{i=1}^{n_0} o_i \in \mathbb{R}^2$, and a scale $s = \sqrt{\frac{1}{n_0} \sum_{i=1}^{n_0} \|o_i - c\|_2^2} + \varepsilon$, is computed, where $\varepsilon > 0$ prevents division by zero. To fix rotation, we compute the principal axis of the centered contour points. Let $O = [o_1 - c; \ldots; o_{n_0} - c] \in \mathbb{R}^{n_0 \times 2}$ and define the covariance

$$\Sigma = \frac{1}{n_0} O^\top O \in \mathbb{R}^{2\times 2}. \tag{27}$$

Let $v \in \mathbb{R}^2$ be the unit eigenvector associated with the largest eigenvalue of Σ, and let $\theta = \operatorname{atan2}(v_2, v_1)$. The canonical rotation is

$$R = \begin{bmatrix} \cos(-\theta) & -\sin(-\theta) \\ \sin(-\theta) & \cos(-\theta) \end{bmatrix}. \tag{28}$$

Since PCA orientation is ambiguous up to sign (v and $-v$), we apply a deterministic sign-fix (e.g., choose the sign such that the skewness of projections onto v is positive), yielding a unique R in practice.

Given any point $p \in \mathbb{R}^2$ (e.g., a hole centroid), its RTS-canonical coordinates are

$$\mathrm{canon}(p) = R\,\frac{(p-c)}{s} \in \mathbb{R}^2. \tag{29}$$

Spatial pyramid Zernike descriptor We implement a spatial pyramid version of the Zernike descriptor retaining Zernike moments' strong global invariance properties while recovering spatial layout information by computing moments not only on the full image but also on a coarse grid of subregions. Let $I \in [0,1]^{\text{height}\times\text{width}}$ denote the normalized grayscale image (after centering/scaling), and let $\mathbf{g} = \{B_i\}_{i=1}^{g}$ be a rectangular grid partition of $[1,H]\times[1,W]$ into $g = g_y g_x$ cells. For any region Re, let $x_{Re}(I)$ denote the corresponding image patch, embedded into a square array by zero-padding and then resampled to a fixed resolution $P \times P$ (so that Zernike polynomials are evaluated on a consistent discrete unit-disk support). Writing $Z_{n,m}(\cdot; r) \in \mathbb{C}$ for the Zernike moment of order (n,m) computed on a disk of radius r (with $n \le d$), we define the level-0 (global) feature vector as $\phi_0(I) = (|Z_{n,m}(I; r0)|)_{(n,m)\in\beth_d}$ and the level-1 (local) feature vectors as $\phi_i(I) = (|Z_{n,m}(\Pi_{B_i}(I); r1)|)_{(n,m)\in\beth_d}$, where $\beth_d = \{(n,m) : 0 \le n \le d, |m| \le n, n-|m| \text{ even}\}$. The final descriptor is the concatenation

$$\phi_{SPZ}(I) = [\phi_0(I); \phi_1(I); \cdots; \phi_{\mathbf{g}}(I)] \in \mathbb{R}^{(\mathbf{g}+1)|\beth_\Gamma|}. \tag{30}$$

Using magnitudes $|Z_{n,m}|$ yields rotation invariance within each region (since rotations induce only a phase shift $e^{-im\alpha}$ in $Z_{n,m}$), while the grid decomposition preserves coarse spatial layout by preventing cancellations across distant parts of the shape.

In order to keep the entire spatial pyramid rotation invariant, we estimate the canonical orientation R from the outer contour. We then rotate the normalized patch into that frame before computing grid features:

$$I_{can}(p) = I(R^{-1}p). \tag{31}$$

Then the grid is effectively attached to the actual object, and concatenating per-cell features becomes approximately rotation invariant.

we implement a group pooling approach over discrete rotations. For a finite rotation group $G_K = \{R_k\}_{k=0}^{K} - 1$ with R_k rotations by $2\pi k/K$, let $\phi(I) \in \mathbb{R}^{G\times M}$ be the per-cell feature tensor (cells x features) or a flattened vector. The rotation invariant max pooling operation is constructed as:

$$\phi_{inv}[j] = \max_k \phi(R_k I)[j]. \tag{32}$$

HOG descriptor While Zernike provides the global structure, it lacks the local structure which introduces subtle variations that differentiates images. Thus, we provide the histogram of oriented gradients (HOG) ϕ_{HOG} features of the images to capture local, orientation-driven stroke structure that is often attenuated by global moment descriptors. HOG represents an image by histograms of local gradient orientations pooled over spatial cells with block normalization, making it well suited to character-like data where discrimination depends on the presence, direction, and relative placement of edges.

HOG is naturally robust to moderate illumination/contrast changes and small geometric perturbations due to its local pooling and normalization. Using both descriptors therefore yields a more informative outer-channel representation: Zernike captures global shape geometry in a low-dimensional, rotation-stable form, while HOG supplies high-frequency, local edge cues, improving separability under noise and partial corruption before the features are mapped into the HD space.

6.1 Proof of Theorem

Proof. The continuous unsupported set $U \subset [0, 2\pi)$ is defined as a set of directions where the boundary localization error exceeds τ:

$$U = \{\theta \in [0,2\pi) : |d_A(\theta) - d_H(\theta)| > \tau\}. \tag{33}$$

The spoke-wheel test samples the continuous domain at N_θ discrete, uniformly spaced angles:

$$\theta_j = \frac{2\pi(j-1)}{N_\theta}, j = 1, \ldots, N_\theta. \tag{34}$$

The angular distance between two adjacent spokes is:

$$\Delta\theta = \frac{2\pi}{N_\theta}. \tag{35}$$

The total number of spokes whose errors are larger should be limited to:

$$N_{unsupported} \leq \frac{Total\ unsupported\ arc\ length}{\Delta\theta} + 2 = N_\theta\epsilon + 2 \tag{36}$$

The unsupported fraction would then be

$$\frac{N_{unsupported}}{N_\theta} \leq \epsilon + \frac{2}{N_\theta}. \tag{37}$$

Thus the enclosure efficiency is given as

$$e_N(H^*) \geq 1 - \epsilon - \frac{2}{N_\theta}. \tag{38}$$

The maximum discrete spoke gap span g_N is given as:

$$g_N(H^*) \leq \frac{N_\theta\gamma + 2}{N_\theta} = \gamma + \frac{2}{N_\theta}. \tag{39}$$

□

PSEUDOCODE

Algorithm 1: Topology-guided HDC with fixed topology compatibility and OnlineHD adaptation.

Input: Training and test sets $(X_{\mathrm{tr}}, y_{\mathrm{tr}})$ and $(X_{\mathrm{te}}, y_{\mathrm{te}})$; validation fraction ρ; corruption $\mathcal{C}$; HV dimension D; classes C; batch size B; OnlineHD epochs E and step size η; hole-extraction parameters $\Theta_{\mathcal{H}}$; alpha grid $\mathcal{A}$; fixed topology weight β.

Output: Predictions before and after OnlineHD, $\widehat{y}_{\mathrm{te}}^{(0)}$ and $\widehat{y}_{\mathrm{te}}^{(1)}$; prototypes (P_h, P_z) and (P'_h, P'_z); selected weights $\alpha^{(0)}$ and $\alpha^{(1)}$.

$(X_{\mathrm{tr}}, y_{\mathrm{tr}}, X_{\mathrm{val}}, y_{\mathrm{val}}) \leftarrow \mathrm{StratifiedSplit}(X_{\mathrm{tr}}, y_{\mathrm{tr}}; \rho)$
$\widetilde{X}_{\mathrm{te}} \leftarrow \mathrm{ApplyCorruption}(X_{\mathrm{te}}; \mathcal{C})$
foreach $S \in \{\mathrm{tr}, \mathrm{val}, \mathrm{te}\}$ **do**
  foreach $x \in X_S$ **do**
    $x' \leftarrow \mathrm{NormalizeGlyph}(x)$
    $f^h(x) \leftarrow \mathrm{HOG}(x')$
    $f^z(x) \leftarrow \mathrm{SpatialPyramidZernike}(x')$
    $\widehat{\mathcal{H}}(x) \leftarrow \mathrm{RobustHoleExtraction}(x'; \Theta_{\mathcal{H}})$
    $z(x) \leftarrow \min\Big(|\widehat{\mathcal{H}}(x)|, 2\Big)$
$\widehat{P}(z \mid c) \leftarrow \mathrm{FitTopologyCompatibility}(\{z(x)\}_{x \in X_{\mathrm{tr}}}, y_{\mathrm{tr}})$
$\mathrm{Enc}_h \leftarrow \mathrm{FitRandomProjection}(\{f^h(x)\}_{x \in X_{\mathrm{tr}}}; D)$
$\mathrm{Enc}_z \leftarrow \mathrm{FitRandomProjection}(\{f^z(x)\}_{x \in X_{\mathrm{tr}}}; D)$
$(P_h, P_z) \leftarrow \mathrm{TrainPrototypes}(X_{\mathrm{tr}}, y_{\mathrm{tr}}; \mathrm{Enc}_h, \mathrm{Enc}_z, B, C)$
$\alpha^{(0)} \leftarrow \mathrm{TuneAlpha}(P_h, P_z, \widehat{P}, \beta; X_{\mathrm{val}}, y_{\mathrm{val}}, \mathcal{A})$
$\widehat{y}_{\mathrm{te}}^{(0)} \leftarrow \mathrm{Predict}(P_h, P_z, \widehat{P}, \alpha^{(0)}, \beta; \widetilde{X}_{\mathrm{te}})$
$(P'_h, P'_z) \leftarrow \mathrm{OnlineHD}(P_h, P_z; X_{\mathrm{tr}}, y_{\mathrm{tr}}, E, \eta, B)$
$\alpha^{(1)} \leftarrow \mathrm{TuneAlpha}(P'_h, P'_z, \widehat{P}, \beta; X_{\mathrm{val}}, y_{\mathrm{val}}, \mathcal{A})$
$\widehat{y}_{\mathrm{te}}^{(1)} \leftarrow \mathrm{Predict}(P'_h, P'_z, \widehat{P}, \alpha^{(1)}, \beta; \widetilde{X}_{\mathrm{te}})$
return $\widehat{y}_{\mathrm{te}}^{(0)}$, $\widehat{y}_{\mathrm{te}}^{(1)}$, (P_h, P_z), (P'_h, P'_z), $\alpha^{(0)}$, $\alpha^{(1)}$

For an image x, prediction uses the fused class score

$$S_c(x; \alpha, \beta) = s_c^{\mathrm{hog}}(x) + \alpha s_c^{\mathrm{zer}}(x) + \beta \log \widehat{P}(z(x) \mid c), \tag{40}$$

where s_c^{hog} and s_c^{zer} are cosine similarities to the corresponding class prototypes. The topology model and β remain fixed during OnlineHD; only the HOG and Zernike prototypes are updated, and α is reselected on the validation set.

HYPERPARAMETER SETTINGS

Table 5: Hyperparameter settings used in the experiments.

Category	**Setting**
HD representation	Dimension $D = 10{,}000$; bipolar HVs in $\{-1, +1\}^D$
Classes	$C = 10$ (MNIST); $C = 26$ (EMNIST Letters)
Data split	Stratified validation fraction $\rho = 0.2$
Batching	Batch size $B = 512$
OnlineHD	$E = 20$ epochs; step size $\eta = 1$
Fusion	$\alpha \in \{0, 0.1, \ldots, 1.0\}$ selected on validation data; topology weight $\beta \in \{0.0, 0.01, 0.025, 0.05, 0.1, 0.2, 0.4\}$ before and after OnlineHD
Random projection	HOG seed 2; Zernike seed 1
Role vectors	HOG seed 1000; Zernike/outer-shape seed 999
Image normalization	Output size 64×64; padding 2 pixels
HOG	9 orientations; 4×4-pixel cells; 2×2-cell blocks; L2-Hys normalization
Zernike	Maximum degree 6; radius 20; one global region and a 2×2 spatial grid
Hole proposals	Morphological closing radii $\mathcal{R} = \{0, 1, 2\}$; minimum area 4 pixels; maximum area $0.3HW$
Spoke validation	$N_\theta = 48$ spokes; ray step 0.5 pixels; enclosure threshold $\eta_e = 0.85$; maximum unsupported angular fraction $\eta_g = 0.15$
Multiscale matching	Centroid threshold 3 pixels; minimum scale count 1
Topology categories	$z(x) \in \{0, 1, 2+\}$
Topology model	Laplace smoothing $\lambda = 1$; fit on clean training data only
Test corruptions	Rotation 20°; Gaussian noise $\sigma \in \{0.1, 0.2\}$; salt-and-pepper probability $p = 0.1$; cutout size $\in \{4, 8, 12\}$; zoom scale $\in \{0.75, 0.5\}$