Hybrid Shape Descriptor Fusion with LightGBM for Robust 3D Mesh Classification and Retrieval

Khadija Arhid[1], Youness Ghazi[2], Ilham Kachbal[3], Fatima Rafii Zakani[4],
Mohcine Bouksim[5], Said El Abdellaoui[6], Taoufiq Gadi[7], Mohamed Aboulfatah[8]
Processes, Signals, Industrial Systems and Computer Science Laboratory,
EST Safi, Cadi Ayyad University, Marrakech, Morocco[1,3,6]
Mathematics, Computer Science and Engineering Sciences Laboratory,
FST, Hassan First University, Settat, Morocco[2,7,8]
Laboratory of Applied Sciences, SOVIA Team, ENSA, Abdelmalek Essaâdi University, Al Hoceïma, Morocco[4]
SmartICT Lab, ENSA, Mohammed First University, Oujda, Morocco[5]

***Abstract*—Recent advances in 3D mesh acquisition and the development of interactive modeling tools have significantly increased both the quantity and diversity of available 3D model databases. Therefore, the task of searching, querying, and retrieving models in large-scale 3D databases has become a focus of research in this area. Indexing 3D models for content-based retrieval is a challenging task that involves numerous algorithms and tools to capture the most significant representation of the object. In this study, a novel framework for 3D mesh retrieval is proposed that combines distribution-based, spectral, and geometric features into a single representation and employs a machine learning classifier based on LightGBM (Light Gradient Boosting Machine) for classifying 3D objects. To capture the complex geometry of 3D meshes, our approach analyzes surface smoothness, radial vertex distributions, spectral signatures, global shape distributions, topological connectivity, and local curvatures. Evaluated on the Princeton Shape Benchmark (PSB), the proposed approach achieves a 1st Tier accuracy of 0.97 and an F-Measure of 0.96, substantially outperforming both individual descriptors and state-of-the-art methods. The mean pairwise cross-correlation between descriptors is low ($\bar{\rho} = 0.128$), confirming their complementary rather than redundant nature. The proposed approach presents a consistent solution with potential applications in various areas, such as computer vision, robotics, e-commerce, medical imaging, and other related fields.**

***Keywords*—3D object; 3D mesh retrieval; classification; machine learning; shape matching; feature fusion; LightGBM**

I. Introduction

The rapid advances in 3D mesh acquisition techniques, combined with significant improvements in interactive modeling software, have greatly expanded and diversified the available databases of 3D models. This evolution has created a strong demand for efficient methods for searching, navigating, and extracting 3D models from large-scale databases [1], [2], [3], [4]. As a result, the classification and retrieval of 3D meshes have become active and important research areas in computer vision and numerous related fields. The main challenge in content-based research is designing descriptors that effectively capture the fundamental geometric properties of 3D meshes, due to their inherent structural complexity [2], [5], [3].

Existing retrieval approaches are based on a spectrum of paradigms, including feature-based descriptors [6], graph-based representations, machine and deep learning approaches [7], [8], [9], [5], and partial shape-based indexing [10]. Despite the wide variety of approaches and the progress made in this field, current methods often struggle to deliver reliable performance in real-world scenarios characterized by highly complex shapes, sensor noise, and incomplete data, which introduce significant variability in real-world conditions.

Single-descriptor approaches face three fundamental failure modes that motivate a hybrid design. First, local descriptors such as the Shape Index capture surface curvature accurately but fail to distinguish objects that share similar local geometry yet differ in global structure, for example, a cup and a vase may have nearly identical curvature profiles but very different global shapes. Second, global descriptors such as D2 lose discriminative power when objects have similar overall shape spread but distinct fine-grained topology, as is common between categories such as Airplane and Bird. Third, spectral descriptors such as HKS degrade under noisy or incomplete meshes, precisely the conditions encountered in real-world 3D acquisition pipelines.

This study presents a framework for 3D mesh retrieval that combines hybrid descriptors with a LightGBM classifier, providing a compact and discriminative description of features. Local distribution, spectral, topological, and geometric descriptors are fused into a single normalized feature vector, providing an accurate and efficient characterization of the geometry of 3D objects.

The remainder of this study is structured as follows: Section II presents existing methods for 3D mesh retrieval. While, Section III describes the proposed method. Section IV presents and analyzes the experimental results. Finally, Section V concludes the study and outlines future research directions.

II. Related Work

3D object classification and retrieval is still a significant challenge in the computer vision community. There are many applications that are still not fully explored such as robotics, medical imaging, cultural heritage and autonomous vehicles systems. Over the past decade, researchers have proposed a

wide variety of methods for extracting meaningful features from 3D meshes to improve classification accuracy. The following sections present an overview of the most relevant approaches proposed in the literature.

A. Geometric Descriptors

Many 3D mesh indexing techniques rely on geometric descriptors to represent 3D models by capturing both global and local shape characteristics. Osada et al. [11] introduced the D2 shape distribution, which employs statistical analysis to represent the global geometry of a 3D mesh by measuring the distribution of pairwise Euclidean distances between randomly sampled surface points. For local surface analysis, Koenderink and van Doorn [12] proposed the shape index, which leverages the minimum and maximum principal curvature values to characterize local surface shape. More recently, transformer-based models [13] have combined machine learning with geometric refinement to improve shape understanding, while Ma et al. [14] introduced a hierarchical architecture that employs Geodesic Distances and curvature at multiple scales for the extraction of detailed shape features.

B. Spectral and Topological Descriptors

The analysis of 3D meshes using eigenvalues and eigenvectors of the Laplace–Beltrami operator, which leads to the study of intrinsic geometry of the meshes using the so-called spectral methods. The Heat Kernel Signature (HKS) proposed by Sun et al. [15] and the Global Point Signature (GPS) proposed by Rustamov [16] are point-wise descriptions of the local and global geometry of a mesh and are widely used in shape retrieval and classification. Topological descriptors derived from the topological data analysis (TDA) of the mesh structure are also used. Persistent homology, derived by Edelsbrunner et al. [17] is the most popular of the topological descriptors and has been shown to be useful in shape analysis by being able to differentiate shapes belonging to the same class but which are geometrically different [18]. In general, however, these approaches are quite expensive, which can severely limit their applicability to very large databases.

C. Machine and Deep Learning Approaches

Machine and deep learning have substantially improved 3D mesh retrieval systems in terms of efficiency, accuracy, and scalability. Recent work employs CatBoost to learn compact descriptors from mesh features [8], [9], while deep learning architectures such as PointNet [19] and VoxelNet [20] learn features directly from raw point clouds and voxel grids, respectively. Bouksim et al. [21] demonstrated effective retrieval using artificial neural networks trained on handcrafted features, and transformer-based models [13], [14] have further improved shape understanding through geometric refinement.

D. Partial Shape-Based Descriptors

Partial indexing methods indexe 3D models based on the decomposition of objects into their constituents and on modeling their signatures from these constituents. Suzuki et al. [22] decomposed polygons into segments using angles between normal vectors of the edges of polygons and used standard descriptors for each segment. Barra et al. [23] proposed a partial shape matching method based on Reeb graph theory. Moumoun et al. [24] combined a D2 [11] descriptor at the segment level with a global shape spectrum descriptor [25]. Our prior work [10] used data envelopment analysis for inter-object comparison and a hybrid similarity score for comparing 3D objects.

Although 3D object classification has been studied in a large number of recent papers, there are still open challenges. Traditional descriptors used in most of the state-of-the-art methods have several limitations: they do not handle well objects from different categories having the same geometric patterns, and are sensitive to noisy or incomplete data. Learning from top of persistent homology and spectral descriptors is highly time-consuming. These limitations motivate the development of hybrid model that preserve a good balance between the classification accuracy and the learning efficiency.

III. Proposed Approach

To overcome the limitations of individual descriptors identified in the preceding review, we propose a multi-descriptor framework that combines complementary shape descriptors into a single normalized feature vector. Relying on a single shape descriptor offers computational simplicity but typically results in weak classification performance due to three fundamental limitations. First, a single descriptor is insufficient to fully represent the geometric complexity of 3D meshes. For example, while the shape index effectively measures local surface curvature, it fails to capture global structural properties. Second, relying on a single descriptor makes the retrieval system unreliable when querying noisy or incomplete databases. While the D2 distribution is highly robust to noise in meshes, its accuracy in incomplete meshes is reduced. Third, describing and distinguishing objects belonging to highly distinct categories but sharing very similar global geometrical shapes is still is a significant challenge for single descriptor approaches. While spectral descriptors fail in providing sufficient discriminant information in the case of objects having highly similar shapes but with different details.

To address these limitations, in this study, we propose a multi-descriptor approach based on several complementary 3D surface shape descriptors. By applying the PCA (Principal Component Analysis) on the concatenated descriptors, we get a unique and normalized vector. Then, this new descriptor is combined with a LightGBM classifier to achieve the desired results for the 3D mesh retrieval task. The proposed pipeline for classification and retrieval of 3D meshes is illustrated in Fig. 1.

A. Theoretical Justification for Descriptor Selection

The design of the proposed multi-descriptor framework is grounded in the principle of complementary feature coverage across four orthogonal representation spaces: local geometry, global geometry, spectral intrinsic properties, and topological structure. Since none of the above descriptors are mutually exclusive for all of the spaces; a principle for selecting a representative descriptor in each space is required. The selection of one such descriptor in each space is sufficient for a unified representation to be effective.

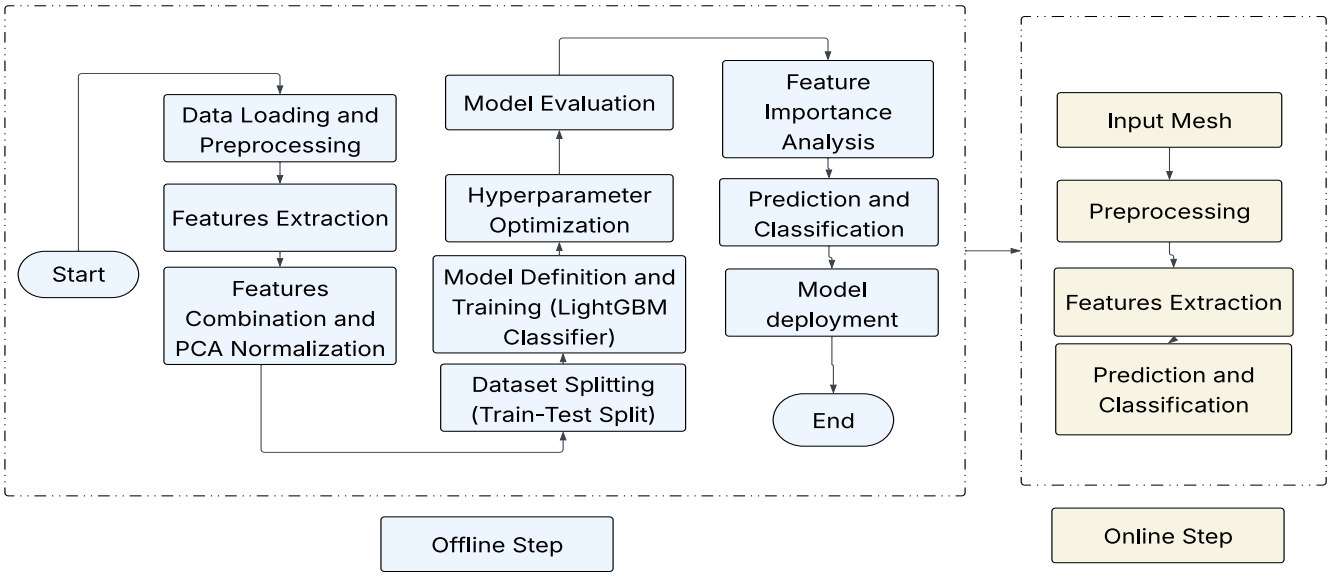

Fig. 1. Sequential steps of the proposed approach.

Formally, let a 3D mesh $\mathcal{M}$ be characterized by a feature space $\mathcal{F} = \{F_{\text{local}}, F_{\text{global}}, F_{\text{spectral}}, F_{\text{topo}}\}$. A descriptor set $\mathcal{D} = \{d_1, \ldots, d_k\}$ achieves complete coverage if and only if, for each subspace $F_s \in \mathcal{F}$, there exists at least one descriptor $d_i \in \mathcal{D}$ that is sensitive to variations in F_s and insensitive to its complement. The five descriptors selected in this work satisfy this coverage criterion, as detailed in Table I.

To validate this selection empirically, we computed pairwise cross-correlation between descriptor histograms on the Princeton Shape Benchmark (Table II) and intra-class consistency for each descriptor (Table III). The Mean inter-descriptor correlation is low ($\bar{\rho} = 0.128$) which means that our optimal descriptors are not heavily redundant. The only large correlation value is for the D2–Geodesic pair ($\rho = 0.777$). This is also expected. The D2 descriptor is computed in Euclidean space whereas the Geodesic Distribution is computed in intrinsic geometry space where distances are computed along the surface following the mesh topology. The average intra-class correlation is $\bar{\rho} = 0.777$ which confirm that each descriptor produces stable, reproducible representations within the same semantic category. our optimal descriptors exhibit the two main properties of feature fusion: Individual reliability in each feature space, and collective non-redundancy in the fused feature space.

B. Pipeline Description

1) Data acquisition and preprocessing: The first step involves loading 3D mesh models from the Princeton Shape Benchmark [26]. Each mesh is normalized by centering it at the origin and scaling it to fit within a unit sphere, thereby ensuring invariance to translation and scale transformations. Let $V = \{v_1, v_2, \ldots, v_N\}$ be the list of N vertices of a 3D mesh, where each vertex $v_i = (x_i, y_i, z_i)$ represents a point in 3D space.

a) Centering (translation tnvariance): Computing the centroid $\mathbf{c}$ (mean position) of all vertices:

$$c = \frac{1}{N} \sum_{i=1}^{N} \mathbf{v}_i \tag{1}$$

Each vertex needs to be translated so that the centroid is moved to the origin:

$$\mathbf{v}'_i = \mathbf{v}_i - \mathbf{c}, \text{ for } i = 1, \ldots, N$$

Let $V' = \{\mathbf{v}'_1, \mathbf{v}'_2, \ldots, \mathbf{v}'_N\}$ denote the set of centered vertices.

b) Scaling (scale invariance): Calculate the maximum Euclidean norm (L_2 distance) from the origin.

$$d_{\max} = \max_{i=1,\ldots,N} \sqrt{(x'_i)^2 + (y'_i)^2 + (z'_i)^2} \tag{2}$$

Each centered vertex is then normalized to ensure scale invariance:

$$\mathbf{v}''_i = \frac{\mathbf{v}'_i}{d_{\max}}, \quad \text{for} \quad i = 1, \ldots, N \tag{3}$$

The resulting set $V'' = \{\mathbf{v}''_1, \mathbf{v}''_2, \ldots, \mathbf{v}''_N\}$ contains the normalized vertices, ensuring that subsequent feature extraction is invariant to the mesh's original position and scale.

2) Multi-descriptor feature extraction: Five complementary descriptors are computed for each normalized mesh to capture multiple geometric and spectral properties. The used descriptors are:

TABLE I. THEORETICAL JUSTIFICATION FOR THE SELECTED DESCRIPTORS

Descriptor	Feature Space	Sensitivity	Limitation Addressed by Others
Shape Index (SI)	Local geometry	Principal curvatures at each vertex	Lacks global context; compensated by D2 and Radius Distribution
D2 Distribution	Global geometry	Pairwise Euclidean distances over the full surface; encodes overall shape spread	Loses local detail; compensated by SI and HKS
Heat Kernel Signature (HKS)	Spectral intrinsic	Multi-scale heat diffusion on the Laplace–Beltrami operator; isometry-invariant	Noise-sensitive; compensated by D2 and Geodesic Distance
Radius Distribution	Global radial structure	Distances from centroid to vertices; encodes symmetry and compactness	Reference-point sensitive; compensated by Geodesic Distance
Geodesic Distance Distribution	Topological surface	Intrinsic surface distances; accounts for holes, handles, and surface curvature	Computationally expensive; subsampled for efficiency

TABLE II. PAIRWISE CROSS-CORRELATION MATRIX BETWEEN DESCRIPTOR HISTOGRAMS.

	SI	D2	HKS	Radius	Geodesic
SI	1.000	-0.141	0.390	-0.269	-0.105
D2	-0.141	1.000	-0.009	0.492	0.777
HKS	0.390	-0.009	1.000	-0.294	0.037
Radius	-0.269	0.492	-0.294	1.000	0.406
Geodesic	-0.105	0.777	0.037	0.406	1.000
Mean $\bar{\rho}$			0.128		
Range			$[-0.294,\ 0.777]$		

TABLE III. INTRA-DESCRIPTOR CONSISTENCY ACROSS THE PRINCETON SHAPE BENCHMARK.

Descriptor	Intra-Consistency (ρ)
Shape Index (SI)	0.852
D2 Distribution	0.841
HKS	0.760
Geodesic Distance	0.763
Radius Distribution	0.671
Mean	**0.777**

a) Shape Index (SI) [12]: Shape index is a local descriptor that describes the intrinsic geometric curvature of a surface at a given point, and gives information relative to the local curvature at each vertex.

$$SI = \frac{2}{\pi} \arctan\left(\frac{k_2 + k_1}{k_2 - k_1 + \epsilon}\right) \tag{4}$$

where, k_1 and k_2 are the principal curvatures at each vertex, and ϵ is a small constant to avoid division by zero.

The shape index gives us the local surface detail and allows us to distinguish between different shapes, on the basis of the characteristics of the local surface geometry. The shape index is very sensitive to noise of the surface and does not give information about the global shape context. In particular, it does not allow us to compare different shapes with the same local distribution of curvature.

b) D2 Shape Distribution [11]: D2 Distribution represents the global distribution of Euclidean distances between pairs of randomly sampled surface points. It computes the distribution of distances between sampled pairs of points and creates a histogram that represents the global shape properties of the object.

$$D_2 = \text{Histogram}\left(\{d(p_i, p_j)\}\right) \tag{5}$$

where, $d(p_i, p_j)$ represents the Euclidean distance between the points p_i and p_j selected randomly on the object's surface.

This global shape descriptor aims at representing global properties of the geometric shape defined at the vertex level, by analyzing the pairwise vertex distances. This approach has the drawback of being very expensive, especially for large models, due to the high number of pairwise vertex distances computed. Also, models with missing data or different mesh resolutions are not handled well in this approach. In particular, the D2 descriptor is highly sensitive to these types of effects and usually loses its discriminative power, especially for quite different objects.

c) Heat Kernel Signature (HKS) [15]: HKS provides a multi-scale spectral signature derived from the Laplace–Beltrami operator that captures both local and global information about the shape of a 3D object. It uses the heat diffusion process on the surface of the 3D model.

$$\text{HKS}(t) = \sum_{i=1}^{N} e^{-\lambda_i t} \phi_i^2 \tag{6}$$

where, λ_i are the eigenvalues of the Laplace–Beltrami operator, ϕ_i are the corresponding eigenfunctions, and t is the diffusion time parameter.

This descriptor describes the local and global geometry of the shape using a heat diffusion process. It is very sensitive to noise and holes in the mesh. Also, computing the eigenvalue decomposition of the Laplacian matrix can be costly, especially

for large meshes. In practice, we compute the HKS on a random subset of the vertices.

d) Radius Distribution [27]: Radius Distribution captures the distribution of distances from mesh vertices to the centroid. It generates a histogram of these distances, which provides insight into the object's overall shape and symmetry.

$$R_{\text{hist}} = \text{Histogram}\left(\{r(p_i)\}\right) \tag{7}$$

where, $r(p_i)$ represents the distance between the reference point and the point p_i on the surface of the 3D object.

This descriptor is easy to compute and gives the object's overall shape. It is very convenient for studying the spatial distribution of points with respect to a center. However, it does not give much information about the local details and so is not very efficient for studying the small-scale geometrical structures of an object. One more drawback is that it is not designed for non-centered objects. Additionally, it is sensitive to the choice of reference point and may lose accuracy if the object is not centered.

e) Geodesic Distance Distribution [17]: Geodesic Distance Distribution captures the distribution of Geodesic Distances between pairs of points on the surface of the 3D object. Unlike Euclidean distances, Geodesic Distances take the surface curvature into account, providing a more accurate representation of the object's shape.

$$G_{\text{hist}} = \text{Histogram}\left(\{d_{\text{geo}}(p_i, p_j)\}\right) \tag{8}$$

where, $d_{\text{geo}}(p_i, p_j)$ is the Geodesic Distance between points p_i and p_j on the object's surface.

This descriptor is more accurate for surface-based 3D mesh analysis as it accounts for the curvature of the object and is effective for distinguishing between shapes with complex surfaces and topologies. However, it is computationally expensive, especially for large meshes, and has sensitivity to the quality and resolution of the mesh.

f) Computational time: The extraction time varies significantly across descriptors, reflecting their different levels of geometric complexity. Shape-based and distribution descriptors are computationally lightweight, while spectral and Geodesic descriptors require more processing due to eigendecomposition of the Laplace–Beltrami operator (HKS) and intrinsic surface distance computation (Geodesic Distance), respectively. Approximate per-mesh extraction times on the PSB are summarized in Table IV.

TABLE IV. APPROXIMATE PER-MESH EXTRACTION TIME

Descriptor	Extraction Time (s)
Radius Distribution	≈ 0.5
Shape Index	≈ 0.8
D2 Distribution	≈ 1.2
HKS	≈ 3.5
Geodesic Distance	$\approx 5.5^*$

*Subsampled for efficiency

3) Feature fusion: After the second step, the histograms generated from the five descriptors are concatenated into a single, high-dimensional feature vector for each 3D model.

4) Dimensionality reduction and normalization: Principal component analysis (PCA) [28], [29] is applied to the combined feature vectors. This step serves both to reduce the dimensionality of the feature space and to normalize the data by decorrelating the features and capturing the principal axes of variation. A specific percentage of variance (95%) is retained. PCA implicitly weights each descriptor by its variance contribution, so descriptors with higher variance contribute more strongly to the retained principal components, providing a data-driven alternative to manual weighting. This normalization also eliminates inter-feature noise and redundancy, producing a more compact and generalizable representation for downstream classification.

5) Dataset splitting: The dataset, now represented by the PCA-transformed feature vectors and corresponding class labels, is divided into distinct training and testing sets using stratified sampling to maintain class proportions.

6) LightGBM classifier training: LightGBM [30] is a high-performance gradient boosting framework designed for optimal speed and memory efficiency. The LightGBM belongs to the family of Gradient Boosting Decision Tree (GBDT) methods, it builds its ensemble by sequentially adding weak decision trees as its core prediction model. The model constructs each new tree to correct the errors made by the previously trained models. The sequential training process enables LightGBM to achieve both high accuracy and efficient computation [31]. LightGBM handles high-dimensional histogram vectors efficiently through histogram-based splits, trains faster than standard GBDT implementations, and its built-in regularization prevents overfitting on the PSB dataset.

7) Hyperparameter optimization: The LightGBM classifier was configured with the following key hyperparameters: learning_rate = 0.06, n_estimators = 2000, min_child_samples = 40, colsample_bytree = 0.7, feature_fraction = 0.7, subsample = 0.7, reg_alpha = 0.1, and reg_lambda = 1.0. Early stopping with a patience of 50 rounds was applied on the validation set to prevent overfitting. The optimal model for the proposed fused descriptor was selected at round 1592, while individual descriptors converged between rounds 103 and 182, confirming the higher representational capacity of the fused feature vector. LightGBM's built-in L1/L2 regularization (`reg_alpha` = 0.1, `reg_lambda` = 1.0) combined with PCA dimensionality reduction effectively constrains model complexity on the high-dimensional fused feature vector.

8) Model training: The LightGBM classifier was trained on the training set using these hyperparameters, with a fixed random state (seed = 42) to ensure reproducibility.

9) Model evaluation: The performance of the trained LightGBM model is evaluated on the held-out test set using multiple complementary metrics: per-class precision and recall, confusion matrix analysis, 1st Tier and 2nd Tier accuracy, Normalized Discounted Cumulative Gain (NDCG), F-Measure, and Precision@K and Recall@K curves.

10)Feature importance analysis: to understand the influence of each principal component derived from the combined

shape descriptors through PCA on the model's predictions, feature importance analysis is performed. This analysis helps us to identify which features most significantly impact the classification performance, offering insights into the underlying geometric characteristics of the 3D meshes.

IV. Results and Discussion

In this section we evaluated the performance of the proposed approach on the Princeton Shape Benchmark (PSB) [26] dataset. The PSB is a 3D model dataset of 2,360 3D models belonging to 19 shape categories. This dataset is one of the standard benchmarks for evaluating the performance of 3D model search engines due to its diversity, and also because it has interesting classes such as humans/armadillos and birds/airplanes which have large numbers of similar 3D models. This diversity and especially this high degree of similarity between many 3D models in the PSB classes make the PSB very challenging for any mesh search algorithm. The evaluation is performed on three orthogonal dimensions. In the first dimension we compared our approach with all the single features in order to measure the gain provided by the fusion of multiple features. In the second dimension we compared our approach with the state-of-the-art methods using a set of evaluation metrics. In the third dimension we performed an ablation study that measures the gain of each individual feature as well as the impact of the classifier and the dimensionality reduction.

To evaluate performance, we employ the following metrics: **Precision** and **Recall** are computed per class. The **Confusion Matrix** summarizes correct and incorrect predictions. **1st Tier** and **2nd Tier** indicate the fraction of queries for which the correct class is ranked first and within the top two, respectively. **NDCG** evaluates ranking quality by considering the positions of relevant items. The **F-Measure** (macro F1) provides a balanced measure across all classes. **Precision@K** and **Recall@K** measure retrieval accuracy and coverage within the top K results.

A. Per-Class Precision and Recall

Fig. 2 and Fig. 3 present the per-class precision and recall scores achieved by each individual descriptor and by the proposed approach across all 19 categories of the Princeton Shape Benchmark.

The results demonstrate the consistent superiority of the proposed approach across all object categories. Each individual descriptor exhibits class-specific weaknesses: the Radius Histogram is sensitive to scale variations and local surface irregularities, leading to performance drops in geometrically complex categories such as Airplane and Ant. The HKS descriptor is sensitive to noise and mesh incompleteness, which reduces its stability across categories. The D2 shape distribution captures global geometry effectively but struggles to discriminate between objects with similar overall shapes and differing fine details. The Shape Index captures precise local curvature information but lacks the global context needed for several object categories.

In contrast, the proposed approach combines all five descriptors into a compact PCA-reduced representation, enabling the LightGBM classifier to generalize more effectively across diverse object types. By leveraging the complementary strengths of each descriptor, the proposed approach achieves near-perfect precision and recall for the majority of categories, substantially outperforming all individual descriptors.

B. Precision@K and Recall@K

Fig. 4 presents the Precision@K and Recall@K curves for all evaluated methods as a function of the number of retrieved shapes K ($K = 1, \ldots, 10$).

As expected, Recall@K increases monotonically with K for all methods. The proposed approach consistently outperforms all individual descriptors at every value of K for both metrics. It achieves very high Precision@1, indicating that the most relevant shape is reliably ranked first, and rapidly approaches Recall@$K = 1.0$, demonstrating that nearly all relevant shapes are retrieved within the top-ranked results.

This superior performance is attributed to the complementary nature of the fused descriptors: HKS captures intrinsic local geometry, the Radius Histogram models the global radial distribution of vertices, the Shape Index characterizes local surface curvature, and the D2 and Geodesic Distance distributions encode global and intrinsic spatial relationships, respectively. Integrating these heterogeneous features allows the proposed approach to exploit the strengths of each individual descriptor while compensating for their respective limitations, yielding a significantly more discriminative and robust similarity measure.

C. Quantitative Evaluation: 1st Tier, 2nd Tier, DCG, and F-Measure

Table V compares the proposed approach against each individual descriptor using four standard retrieval and classification metrics: 1st Tier, 2nd Tier, Discounted Cumulative Gain (DCG), and F-Measure. The proposed approach achieves the highest scores on all four metrics, with a 1st Tier of 0.97, a 2nd Tier of 0.96, a DCG of 0.97, and an F-Measure of 0.96. This represents a substantial improvement over the best-performing individual descriptor, the Shape Index (1st Tier: 0.77, F-Measure: 0.66), confirming the significant advantage of multi-descriptor fusion over any single descriptor.

This improvement is primarily attributed to the integration of multiple complementary descriptors, which enables the model to capture a richer and more discriminative representation of 3D meshes. By combining geometric, topological, and spectral features, the method effectively addresses the limitations of individual descriptors, yielding enhanced robustness against intra-class variability and complex shape structures.

TABLE V. Comparison of the Proposed Approach and Individual Descriptors.

Method	1st Tier	2nd Tier	DCG	F-Measure
Proposed Approach	**0.97**	**0.96**	**0.97**	**0.96**
D2 Distribution	0.70	0.51	0.75	0.50
Shape Index	0.77	0.66	0.82	0.66
HKS	0.58	0.43	0.69	0.40
Radius Histogram	0.73	0.56	0.78	0.52
Geodesic Distance	0.71	0.52	0.75	0.49

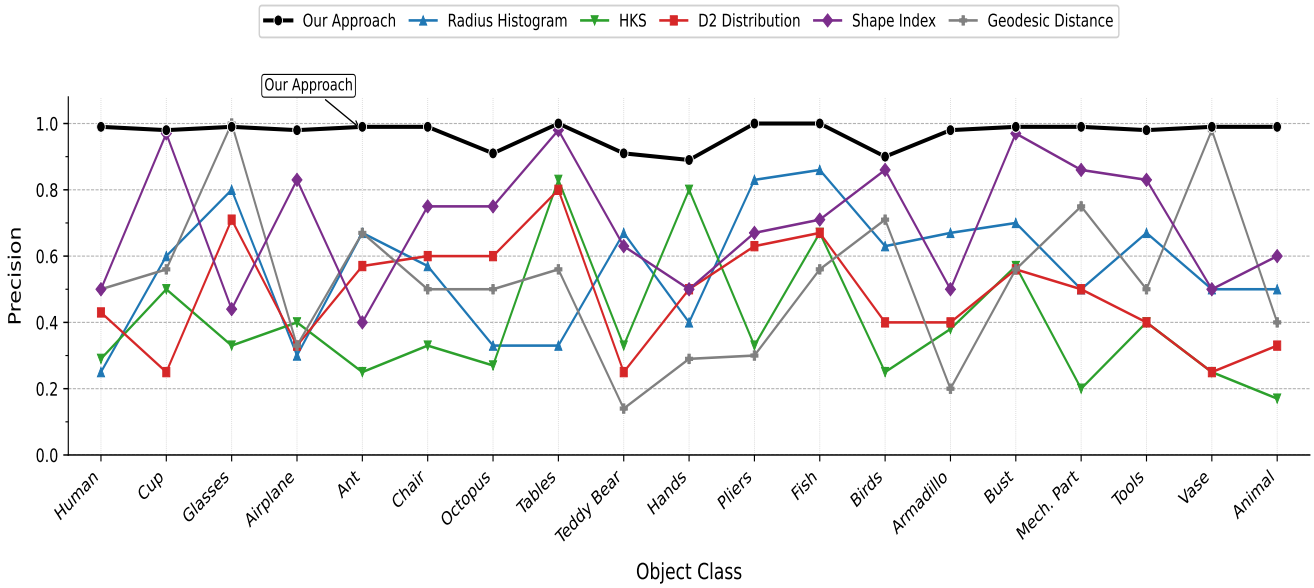

Fig. 2. Per-class precision scores on the Princeton Shape Benchmark.

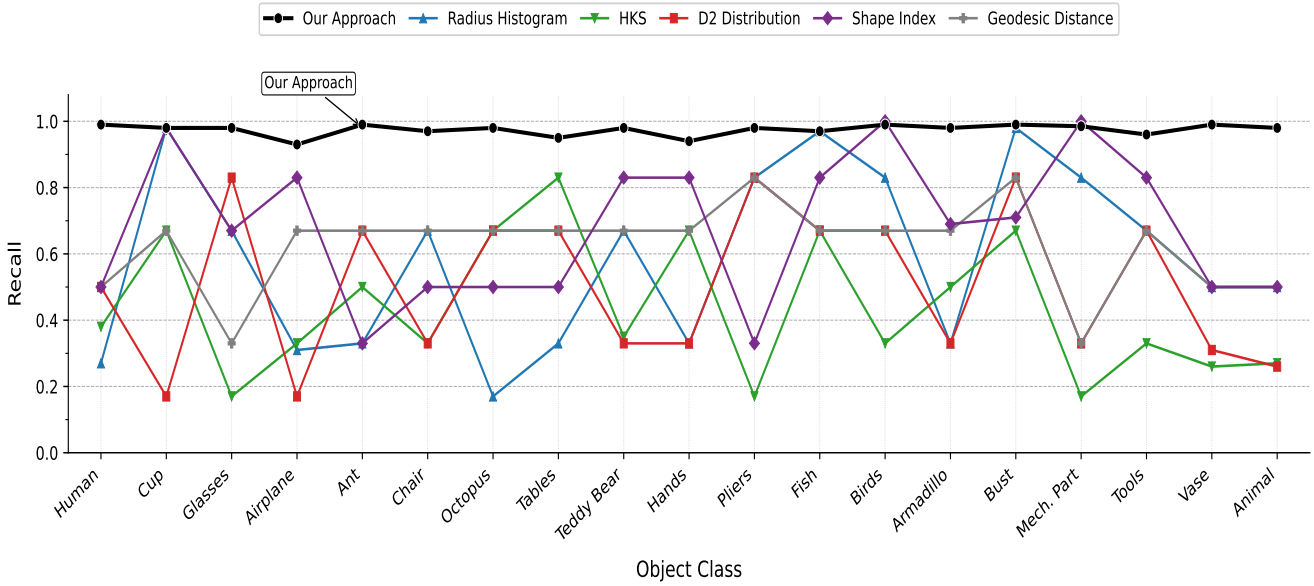

Fig. 3. Per-class recall scores on the Princeton Shape Benchmark.

Table VI provides a comprehensive comparison between the proposed approach and Panorama, LightField, DEA, and ANN across multiple retrieval and classification metrics. The proposed approach achieves outstanding performance, with scores of 0.97 for 1st Tier, DCG, and F-Measure, and 0.96 for 2nd Tier, excelling in both top-ranked retrieval accuracy and overall ranking quality. In contrast, the compared methods yield considerably lower scores, particularly in 2nd Tier and F-Measure. For example, DEA and LightField achieve F-Measure values below 0.40, and even the best-performing competitor, ANN, achieves only 0.80 in 1st Tier and 0.45 in 2nd Tier — substantially below the proposed approach. This advantage is attributed to the advanced feature representation obtained by fusing multiple complementary descriptors, which improves robustness to intra-class variability, noise, and partial occlusions, confirming the framework's suitability for practical 3D mesh analysis applications.

D. Confusion Matrix Analysis

The matrices in Fig. 5 confirm the quantitative superiority of the proposed approach. As demonstrated in Fig. 5a, our proposed descriptor has a diagonal that is almost perfect and very faint off-diagonal entries. In contrast, the single descriptor matrices have much weaker diagonals and larger off-diagonal scatter. Among the single descriptors, HKS has the largest scatter, indicating that it is the least discriminative among the

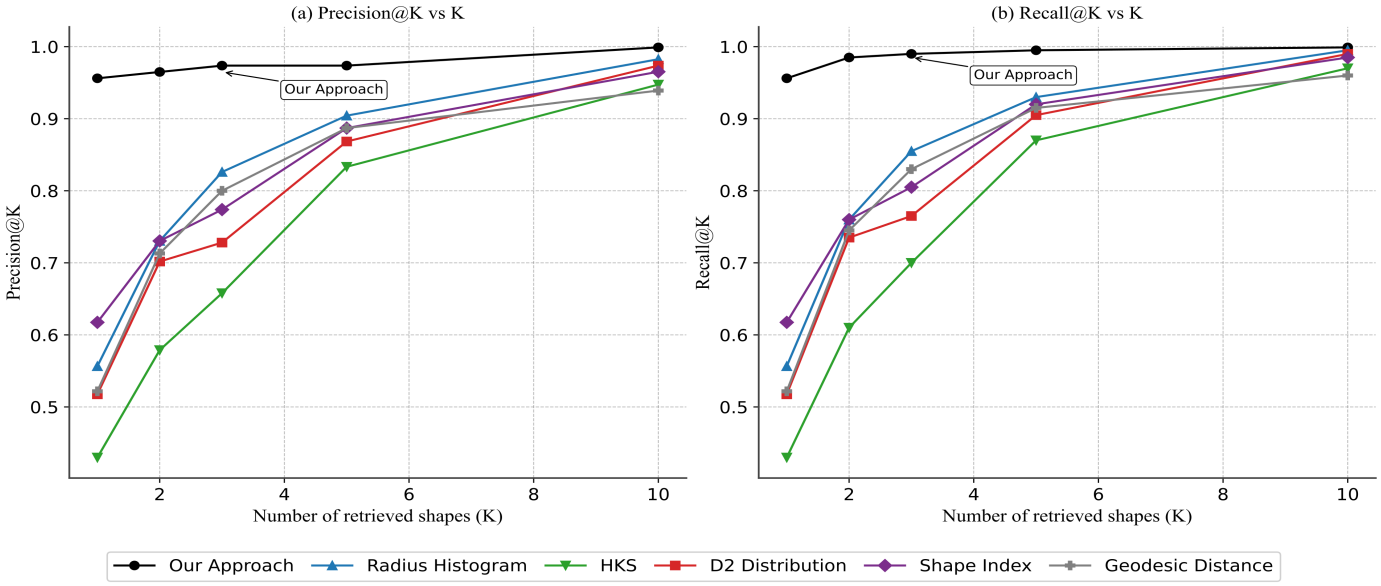

Fig. 4. Precision@K and Recall@K curves for all evaluated methods.

TABLE VI. COMPARISON OF THE PROPOSED APPROACH WITH STATE-OF-THE-ART METHODS

Method	1st Tier	2nd Tier	DCG	F-Measure
Proposed Approach	**0.97**	**0.96**	**0.97**	**0.96**
Panorama	0.73	0.43	0.92	0.42
LightField	0.57	0.36	0.86	0.38
DEA	0.53	0.35	0.82	0.36
ANN	0.80	0.45	0.86	0.38

single descriptors. Although the diagonal of Shape Index is the strongest, it has a larger scatter than our proposed descriptors and the discriminability of the classifier is largely improved by using the combined descriptors. By complementing the curvature, distribution, spectral and distance properties of the shape descriptors the proposed approach effectively get a very discriminative and robust classification model.

E. Ablation Study

To rigorously evaluate the individual contribution of each component within the proposed framework, a systematic ablation study was carried out across two complementary dimensions: the incremental effect of each shape descriptor on classification and retrieval performance, and the influence of classifier selection and dimensionality reduction strategies on the overall system behavior.

1) Descriptor ablation: Incremental addition: Table VII reports the retrieval and classification performance of individual descriptors and their combinations under an incremental feature addition protocol. The baseline is given by the Shape Index and is then successively incremented by adding the other descriptors one after the other in cumulative fashion. The results reveal a consistent and monotonic improvement across all evaluation metrics as each descriptor is introduced. Starting from the Shape Index alone, which attains an F-Measure of 0.66, each successive addition yields a measurable performance gain, culminating in an F-Measure of 0.96 for the full five-descriptor combination. The incorporation of the Geodesic Distance Distribution contributes the largest single increment ($\Delta F_1 = +0.12$), consistent with its capacity to encode intrinsic, surface-aware geometric information that remains inaccessible to Euclidean-based descriptors. This gain is particularly pronounced for topologically complex categories such as Armadillo and Animal, where surface curvature and connectivity play a decisive role in shape discrimination.

TABLE VII. INCREMENTAL DESCRIPTOR COMBINATION RESULTS ON THE PRINCETON SHAPE BENCHMARK.

Configuration	1st Tier	2nd Tier	DCG	F-Meas.	ΔF1
SI only	0.77	0.66	0.82	0.66	—
SI + D2	0.82	0.71	0.86	0.71	+0.05
SI + D2 + HKS	0.86	0.77	0.89	0.77	+0.06
SI + D2 + HKS + Rad.	0.91	0.84	0.93	0.84	+0.07
Full (all 5)	**0.97**	**0.96**	**0.97**	**0.96**	**+0.12**

Rad. = Radius Histogram; Geo. = Geodesic Distance.

2) Component ablation: Classifier and dimensionality reduction: Table VIII reports the performance of the proposed pipeline under varying classifier and dimensionality reduction configurations, evaluated on the full five-descriptor feature vector. Two principal findings emerge from this analysis. First, applying PCA at 95% variance retention yields a higher F-Measure (0.97) than retaining the raw concatenated feature vector without dimensionality reduction (0.94), despite the reduction in feature space dimensionality. This outcome confirms that PCA does not discard discriminative information; rather, it eliminates noise and inter-feature redundancy, producing a more compact and generalizable representation for downstream classification. Retaining 99% variance offers no

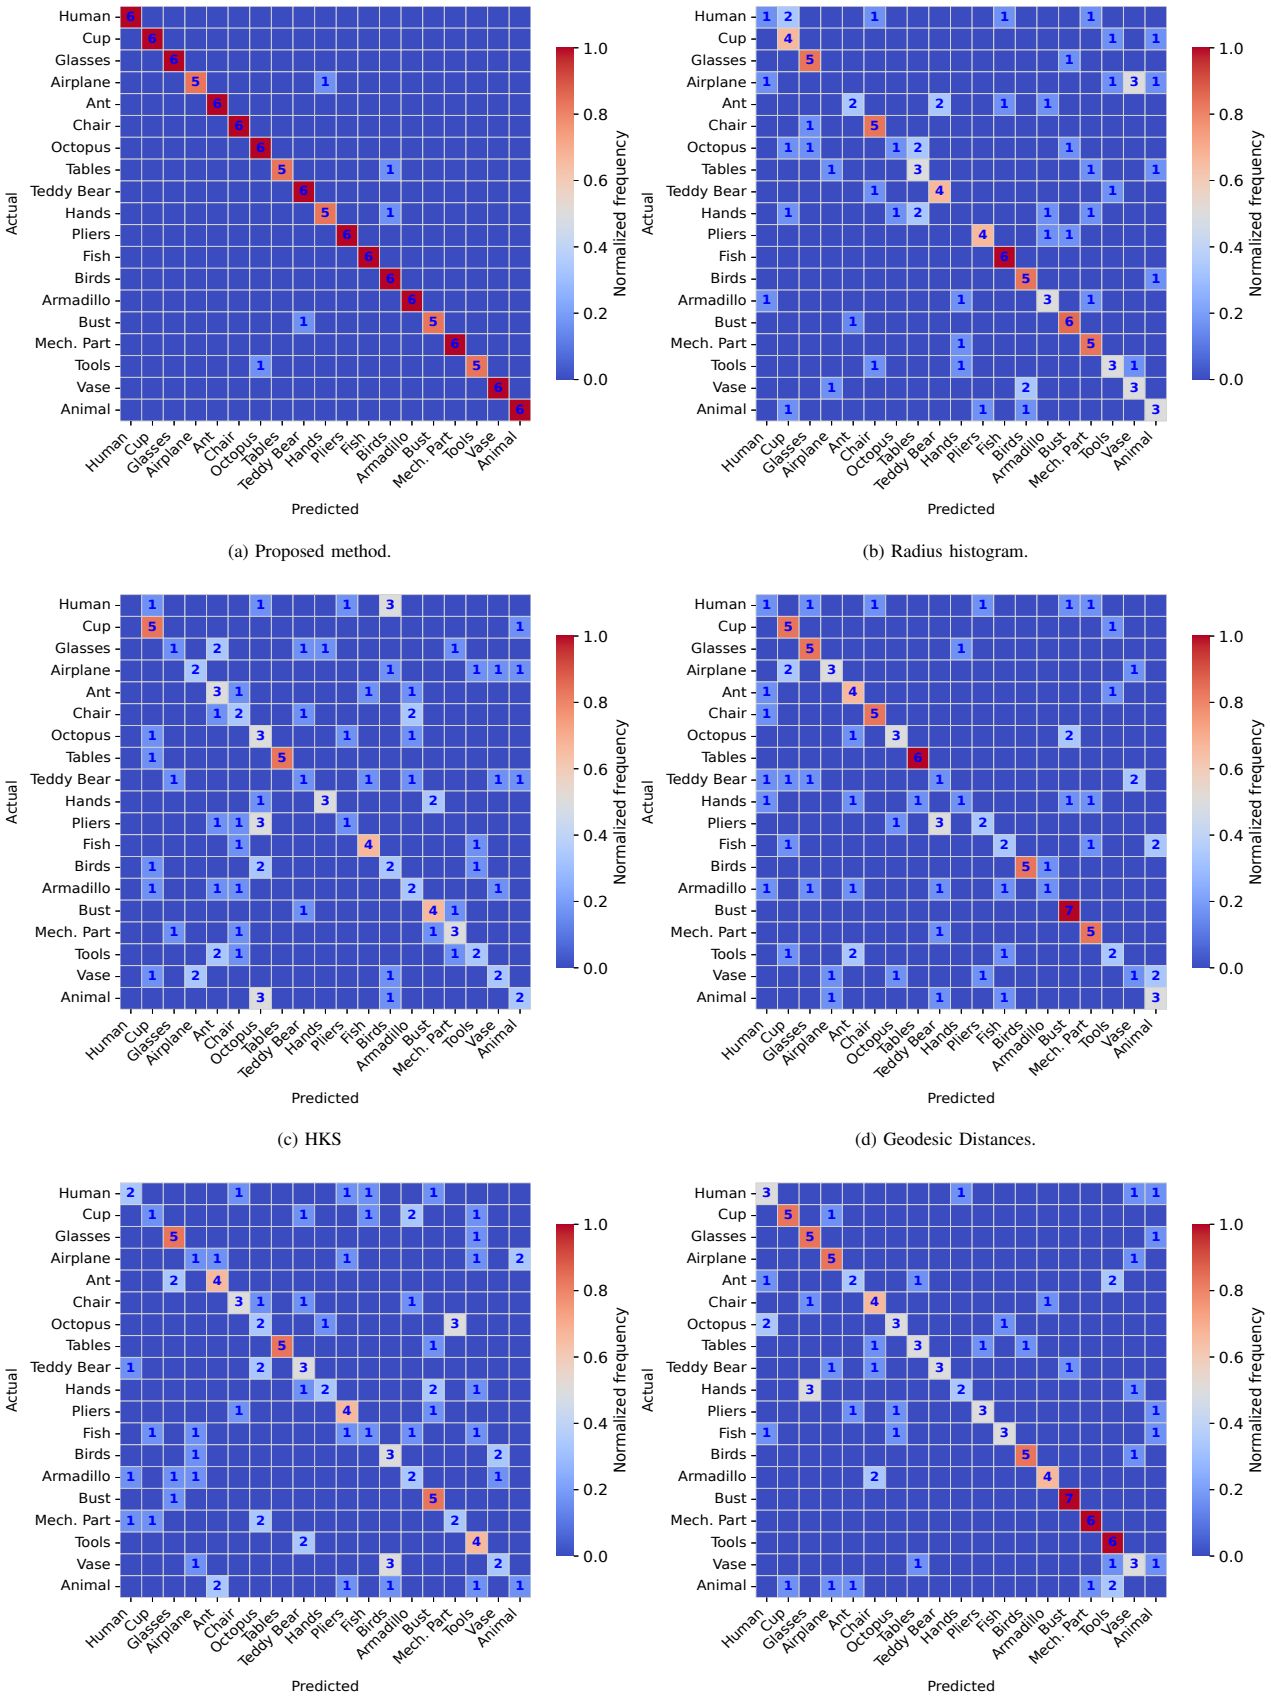

(a) Proposed method.

(b) Radius histogram.

(c) HKS

(d) Geodesic Distances.

(e) D2 Distribution.

(f) Shape index.

Fig. 5. Confusion matrices of the evaluated descriptors.

measurable improvement over the 95% threshold, suggesting that the marginal variance captured beyond this point corresponds to noise rather than meaningful geometric signal. Second, LightGBM consistently outperforms both the Support Vector Machine (SVM) and Random Forest alternatives across all evaluation metrics. This advantage confirms the suitability of LightGBM as the classification backbone of the proposed framework, particularly in scenarios involving high-dimensional, heterogeneous feature vectors derived from multiple complementary descriptors.

TABLE VIII. COMPONENT ABLATION: EFFECT OF CLASSIFIER AND DIMENSIONALITY REDUCTION ON RETRIEVAL AND CLASSIFICATION PERFORMANCE.

Configuration	1st Tier	F-Measure
No PCA, LightGBM	0.94	0.93
PCA (95%), SVM	0.89	0.87
PCA (95%), Rand. Forest	0.92	0.91
PCA (99%), LightGBM	0.96	0.95
PCA (95%), LightGBM (proposed)	**0.97**	**0.96**

All configurations use the full five-descriptor feature vector.

V. CONCLUSION

This study presented a multi-descriptor framework for 3D mesh classification and retrieval combining five complementary shape descriptors, covering local geometry, global structure, spectral intrinsic properties, and topological surface characteristics — into a single PCA-normalized feature vector classified by LightGBM. Two key technical lessons emerge: 1) low inter-descriptor correlation ($\bar{\rho} = 0.128$) is a necessary condition for effective fusion, ensuring that each descriptor contributes non-redundant information; and 2) PCA at 95% variance retention outperforms both raw concatenation and higher retention thresholds by eliminating inter-feature noise rather than discarding discriminative signal. Experimental results on the Princeton Shape Benchmark achieve a 1st Tier of 0.97 and F-Measure of 0.96, outperforming both individual descriptors and state-of-the-art methods. Current limitations include sensitivity to very sparse or heavily corrupted meshes, difficulty in categories with extreme intra-class shape variability such as Animal, and Geodesic computation overhead at scales significantly larger than PSB. Future work will investigate deep learning-based descriptors, robustness experiments on corrupted meshes, and adaptive descriptor weighting conditioned on category-level uncertainty.

REFERENCES

[1] T. Funkhouser, M. Kazhdan, P. Min, and P. Shilane, "Shape-based retrieval and analysis of 3d models," *Communications of the ACM*, vol. 48, no. 6, pp. 58–64, 2005.

[2] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "Shapenet: An information-rich 3d model repository," 2015, arXiv preprint arXiv:1512.03012.

[3] A. Lamba and R. Bhalla, "A review of various 3-d modelling techniques and an introduction to point clouds," in *Proceedings of the International Conference on Emerging Trends in Artificial Intelligence and Smart Systems (THEETAS 2022)*. Jabalpur, India: European Alliance for Innovation, 2022.

[4] A. Martadiansyah, H. K. Putra, M. R. Ramadhan, Ermatita, Abdiansah, and Erwin, "3d fusion hierarchical net reconstruction from 2d transcerebellar images with deep learning," *Engineering Letters*, vol. 32, no. 4, pp. 701–712, 2024.

[5] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun, "Deep learning for 3d point clouds: A survey," 2019, arXiv preprint arXiv:1912.12033.

[6] K. Hussein and A. Behrad, "Three-dimensional shape recognition and classification using local features of model views and sparse representation of shape descriptors," *Journal of Information Processing Systems*, vol. 16, no. 2, pp. 343–359, 2020.

[7] M. Gao, N. Ruan, J. Shi, and W. Zhou, "Deep neural network for 3d shape classification based on mesh feature," *Sensors*, vol. 22, no. 18, p. 7040, 2022.

[8] M. Bouksim, F. R. Zakani, K. Arhid, A. Dahbi, T. Gadi, and M. Aboulfatah, "New 3d shape descriptor extraction using catboost classifier for accurate 3d model retrieval," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 5, 2024.

[9] M. Bouksim, M. A. Madani, F. R. Zakani, K. Arhid, T. Gadi, and M. Aboulfatah, "Catboost3d: A novel catboost-based approach for efficient classification of 3d models," in *2024 International Conference on Computing, Internet of Things and Microwave Systems (ICCIMS)*, 2024, pp. 1–5.

[10] K. Arhid, F. R. Zakani, B. Sirbal, M. Bouksim, M. Aboulfatah, and T. Gadi, "A novel approach for partial shape matching and similarity based on data envelopment analysis," *Computer Optics*, vol. 43, no. 2, pp. 316–323, 2019.

[11] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," *ACM Transactions on Graphics*, vol. 21, no. 4, pp. 807–832, 2002.

[12] J. J. Koenderink and A. J. van Doorn, "Surface shape and curvature scales," *Image and Vision Computing*, vol. 10, no. 8, pp. 557–564, 1992.

[13] J. Zhao, L. Jiao, C. Wang, X. Liu, F. Liu, L. Li, and S. Yang, "Geoformer: A geometric representation transformer for change detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–17, 2023.

[14] J. Ma, X. Jiang, A. Fan, J. Jiang, and J. Yan, "Image matching from handcrafted to deep features: A survey," *International Journal of Computer Vision*, vol. 129, no. 1, pp. 23–79, 2021.

[15] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," *Computer Graphics Forum*, vol. 28, no. 5, pp. 1383–1392, 2009.

[16] R. M. Rustamov, "Laplace-beltrami eigenfunctions for deformation invariant shape representation," in *Proceedings of the Eurographics Symposium on Geometry Processing (SGP 2007)*. Eurographics Association, 2007, pp. 225–233.

[17] H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," *Discrete and Computational Geometry*, vol. 28, no. 4, pp. 511–533, 2002.

[18] C. Hofer, R. Kwitt, M. Niethammer, and A. Uhl, "Deep learning with topological signatures," 2017, arXiv preprint arXiv:1707.04041.

[19] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," 2016, arXiv preprint arXiv:1612.00593.

[20] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4490–4499.

[21] M. Bouksim, K. Arhid, F. R. Zakani, M. Aboulfatah, and T. Gadi, "New approach for 3d mesh retrieval using artificial neural network and histogram of features," *Scientific Visualization*, vol. 10, no. 2, pp. 84–94, 2018.

[22] M. T. Suzuki, Y. Yaginuma, T. Yamada, and Y. Shimizu, "A partial shape matching method for 3d model databases," in *IEEE International Conference on Shape Modeling and Applications (SMI)*. IEEE Computer Society, 2009, pp. 33–40.

[23] V. Barra and S. Biasotti, "3d shape retrieval using kernels on extended reeb graphs," *Pattern Recognition*, vol. 46, no. 11, pp. 2985–2999, 2013.

[24] L. Moumoun, M. Chahhou, M. E. Far, A. Haqiq, and T. Gadi, "3d object retrieval using a global-partial analogy and the bayesian approach," in *Proceedings of the 7th International Conference on Signal Image Technology and Internet-Based Systems (SITIS 2011)*, 2011, pp. 502–507.

[25] T. Zaharia and F. Prêteux, "3d shape-based retrieval within the mpeg-7 framework," Institut National des Télécommunications, Tech. Rep., 2001.

[26] X. Chen, A. Golovinskiy, and T. Funkhouser, "A benchmark for 3d mesh segmentation," *ACM Transactions on Graphics*, vol. 28, no. 3, pp. 1–12, 2009.

[27] A. Khodakovsky, O. Sorkine, and M. Gross, "Bounds on approximating 3d shapes with spheres," *Computer Graphics Forum*, vol. 22, no. 3, pp. 30–38, 2003.

[28] R. Bro and A. K. Smilde, "Principal component analysis," *Analytical Methods*, vol. 6, no. 9, pp. 2812–2831, 2014.

[29] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, vol. 61, no. 3, pp. 611–622, 1999.

[30] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," in *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, vol. 30. Curran Associates, Inc., 2017, pp. 3146–3154.

[31] M. Hajihosseinlou, A. Maghsoudi, and R. Ghezelbash, "A novel scheme for mapping of mvt-type pb-zn prospectivity: Lightgbm, a highly efficient gradient boosting decision tree machine learning algorithm," *Natural Resources Research*, vol. 32, no. 6, pp. 2417–2438, 2023.